

Regret bounds for Narendra-Shapiro bandit algorithms

Sébastien Gadat*, Fabien Panloup†, Sofiane Saadane‡

January 19, 2016

Abstract

Narendra-Shapiro (NS) algorithms are bandit-type algorithms developed in the 1960s which have been deeply studied in infinite horizon but for which scarce non-asymptotic results exist. In this paper, we focus on a non-asymptotic study of the *regret* and address the following question: are Narendra-Shapiro bandit algorithms competitive from this point of view? In our main result, we obtain some uniform explicit bounds for the regret of *(over)-penalized-NS* algorithms.

We also extend to the multi-armed case some convergence properties of penalized-NS algorithms towards a stationary Piecewise Deterministic Markov Process (PDMP). Finally, we establish some new sharp mixing bounds for these processes.

Keywords: Regret, Stochastic Bandit Algorithms, Piecewise Deterministic Markov Processes

1 Introduction

The so-called Narendra-Shapiro bandit algorithm (referred to as NSa) was introduced in [19] and developed in [18] as a linear learning automata. This algorithm has been primarily considered by the probabilistic community as an interesting benchmark of stochastic algorithm. More precisely, NSa is an example of recursive (non-homogeneous) Markovian algorithm, topic whose almost complete historical overview may be found in the seminal contributions of [11] and [13].

NSa belongs to the large class of bandit-type policies whose principle may be sketched as follows: a d -armed bandit algorithm is a procedure designed to determine which one, among d sources, is the most profitable without spending too much time on the wrong ones. In the simplest case, the sources (or arms) randomly provide some rewards whose values belong to $\{0; 1\}$ with Bernoulli laws. The associated probabilities of success (p_1, \dots, p_d) are unknown to the player and his goal is to determine the most efficient source, *i.e.* the highest probability of success.

Let us now remind a rigorous definition of admissible sequential policies. We consider d independent sequences $(A_n^i)_{n \geq 0}$ of *i.i.d.* Bernoulli random variables $\mathcal{B}(p_i)$. Each A_n^i represents the reward associated with the arm i at time n . We then consider some sequential predictions where at each stage n a forecaster chooses an arm I_n , receives a reward $A_n^{I_n}$ and then uses this information to choose the next arm at step $n + 1$. As introduced in the pioneering work [20], the rewards are sampled independently of a fixed product distribution at each step n . The innovations here at time n are provided by $(I_n, A_n^{I_n})$ and we are naturally led to introduce the filtration $(\mathcal{F}_n)_{n \geq 0} := \left(\sigma((I_1, A_1^{I_1}), \dots, (I_n, A_n^{I_n})) \right)_{n \geq 0}$. In the following, the sequential admissible policies will be a $(\mathcal{F}_n)_{n \geq 0}$ (inhomogeneous) Markov chain. We also define another filtration by adding all the events before step n and observe that $(\bar{\mathcal{F}}_n)_{n \geq 0} := (\sigma((I_1, (A_1^j)_{1 \leq j \leq d}), \dots, (I_n, (A_n^j)_{1 \leq j \leq d})))_{n \geq 0}$.

*Toulouse School of Economics, UMR 5604, Université Toulouse I Capitole, France. E-mail: sebastien.gadat@math.univ-toulouse.fr.

†Institut de Mathématiques de Toulouse, UMR 5219, 118, route de Narbonne F-31062 Toulouse Cedex 9, France. E-mail: fabien.panloup@math.univ-toulouse.fr

‡Institut de Mathématiques de Toulouse, UMR 5219, 118, route de Narbonne F-31062 Toulouse Cedex 9, France. E-mail: sofiane.saadane@math.univ-toulouse.fr

To sum-up, $\bar{\mathcal{F}}_n$ contains all the results of each arm between time 1 and n although \mathcal{F}_n only provides partial information about the tested arms.

In this paper, we focus on the stochastic NSa whose principle is very simple: it consists in sampling one arm according to a probability distribution on $\{1, \dots, d\}$, and in modifying this probability distribution in terms of the reward obtained with the chosen arm. From this point of view, this algorithm bears similarities with the EXP3 algorithm (and many of its variants) introduced in [5]. Among other close bandit algorithms, one can also cite the Thompson Sampling strategy where the random selection of the arm is based on a Bayesian posterior which is updated after each result. We refer to [1] for a recent theoretical contribution on this algorithm.

Instead of sampling one arm sequentially according to a randomized decision, other algorithms define their policy through a deterministic maximization procedure at each iteration. Among them, we can mention the UCB algorithm [4] and its derivatives (including MOSS [2] and KL-UCB [9]), whose dynamics are dictated by an appropriate empirical upper confidence bound of the estimated best performance.

Let us now present the NSa algorithm. In fact, we will distinguish two types of NSa: crude-NSa and penalized-NSa. Before going further, let us recall their mechanism in the case of $d = 2$ (the general case will be introduced in Section 2). Designating X_n as the probability of drawing arm 1 at step n and $(\gamma_n)_{n \geq 0}$ as a decreasing sequence of positive numbers that tends to 0 when n goes to infinity, crude-NS is recursively defined by:

$$X_{n+1} = X_n + \begin{cases} \gamma_{n+1}(1 - X_n) & \text{if arm 1 is selected and wins} \\ -\gamma_{n+1}X_n & \text{if another arm is selected and wins} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Note that the construction is certainly symmetric, *i.e.*, $1 - X_n$ (which corresponds to the probability of drawing arm 2) has a symmetric dynamics. The long-time behavior of some NSa was extensively investigated in the last decade. To name a few, in [17] and [15], some convergence and rate of convergence results are proved. However, these results strongly depend on both (γ_n) and the probabilities of success of the arms. In order to get rid of these constraints, the authors then introduced in [16] a penalized NSa and proved that this method is an efficient distribution-free procedure, meaning that it unconditionally converges to the best arm on the unknown probabilities p_1 and p_2 . The idea of the penalized-NS algorithm is to also take the failures of the player into account and to reduce the probability of drawing the tested arm when it loses. Designating $(\rho_n)_{n \geq 0}$ as a second positive sequence, the dynamics of the penalized NSa is given by :

$$X_{n+1} = X_n + \begin{cases} \gamma_{n+1}(1 - X_n) & \text{if arm 1 is selected and wins} \\ -\gamma_{n+1}X_n & \text{if arm 2 is selected and wins} \\ -\rho_n\gamma_{n+1}X_n & \text{if arm 1 is selected and loses} \\ \rho_{n+1}\gamma_{n+1}(1 - X_n) & \text{if arm 2 is selected and loses.} \end{cases} \quad (2)$$

Performances of bandit algorithms. In view of potential applications, it is certainly important to have some informations about the performances of the used policies. To this end, one first needs to define what is a “good” sequential algorithm. The primary efficiency requirement is the ability of the algorithm to asymptotically recover the best arm. In [16], this property is referred to as the *infallibility* of the algorithm. If without loss of generality, the first arm is assumed to be the best, (*i.e.* that $p_1 > \max\{p_2, \dots, p_d\}$) and if $X_n^{(1)}$ denotes the probability of drawing arm 1, the algorithm is said to be infallible if

$$\mathbb{P}(X_n^{(1)} \xrightarrow{n \rightarrow +\infty} 1) = 1. \quad (3)$$

An alternative way for describing the efficiency of a method is to consider the behaviour of the cumulative reward S_n obtained between time 1 and n :

$$S_n := \sum_{k=1}^n A_k^{I_k}.$$

In particular, in the old paper [20], Robbins is looking for algorithms such that

$$p_1 - \frac{\mathbb{E}[S_n]}{n} \xrightarrow{n \rightarrow +\infty} 0.$$

This last property is weaker than the infallibility of an algorithm since the Lebesgue theorem associated to (3) implies the convergence above.

A much stronger requirement involves the *regret* of the algorithm. The regret measures the gap between the cumulative reward of the best player and the one induced by the policy. The regret R_n is the $\bar{\mathcal{F}}_n$ -measurable random variable defined as:

$$R_n := \max_{1 \leq j \leq d} \sum_{k=1}^n [A_k^j - A_k^{I_k}]. \quad (4)$$

A good strategy corresponds to a selection procedure that minimizes the expected regret $\mathbb{E}R_n$, optimal ones being referred to as *minimax* strategies.

The former expected regret cannot be easily handled and is generally replaced in statistical analysis by the *pseudo-regret* defined as

$$\bar{R}_n := \max_{1 \leq j \leq d} \mathbb{E} \sum_{k=1}^n [A_k^j - A_k^{I_k}]. \quad (5)$$

Since $p_1 > p_j, \forall j \neq 1$, \bar{R}_n can also be written as

$$\bar{R}_n = \sum_{k=1}^n \mathbb{E}(A_k^1) - \mathbb{E} \left(\sum_{k=1}^n A_k^{I_k} \right) = n \left(p_1 - \frac{\mathbb{E}[S_n]}{n} \right).$$

A low *pseudo-regret* property then means that the quantity

$$n \left(p_1 - \frac{\mathbb{E}[S_n]}{n} \right)$$

has to be small, in particular sub-linear with n . The quantities R_n and \bar{R}_n are closely related and it is reasonable to study the pseudo-regret instead of the true regret, owing to the next proposition:

Proposition 1.1. (i) For any $(\mathcal{F}_n)_{n \geq 0}$ -measurable strategy, we obtain after n plays:

$$0 \leq \mathbb{E}R_n - \bar{R}_n \leq \sqrt{\frac{n \log d}{2}}.$$

(ii) Furthermore, for every integer n and d and for any (admissible) strategy,

$$\sup_{p_1 > p_2 \geq \dots \geq p_d} \mathbb{E}[R_n] \geq \frac{1}{20} \sqrt{nd}.$$

We refer to Proposition 34 of [3] for a detailed proof of (i) and to Theorem 5.1 of [5] for (ii). As mentioned in (ii), the bounds are distribution-free (uniform in p).¹ Since the MOSS method of [2] satisfies $\bar{R}_n \leq 25\sqrt{nd}$, (i) and (ii) show that a non-asymptotic distribution-free minimax rate is on the order of \sqrt{n} .

In particular, a fallible algorithm (meaning that $\mathbb{P}(X_n \xrightarrow{n \rightarrow +\infty} 1) < 1$) necessarily generates a linear regret and is not optimal. For example, in the case $d = 2$, the dependence of \bar{R}_n in terms of (X_n) is as follows:

$$\begin{aligned} \bar{R}_n &= p_1 n - \sum_{k=1}^n (p_1 \mathbb{E}[X_k] + p_2 \mathbb{E}[1 - X_k]) = (p_1 - p_2) \mathbb{E} \left[\sum_{k=1}^n (1 - X_k) \right] \\ &\gtrsim (p_1 - p_2) \mathbb{P}(X_\infty = 0) \times n. \end{aligned} \quad (6)$$

¹The rate orders are strongly different if a dependence in p is allowed.

Objectives. In this paper, we therefore propose to focus on the regret and to answer to the question “Are NSa competitive from a regret viewpoint? In the case of positive answer, what are the associated upper-bounds ?”

Due to some too restrictive conditions of infallibility, it will be seen that the crude-NSa cannot be competitive from a regret point of view. As mentioned before, the penalized NSa is more robust and is *a priori* more appropriate for this problem. More precisely, the penalty induces more balance between exploration and exploitation, *i.e.* between playing the best arm (the one in terms of the past actions) and exploring new options (playing the suboptimal arms). In this paper, we are going to prove that, up to a slight reinforcement, it is possible to obtain some competitive bounds for the regret of this procedure. The slightly modified penalized algorithm will be referred to as the *over-penalized-algorithm* below.

Outline. The paper is organized as follows : Section 2.1 provides some basic information about the crude NSa. Then, in Section 2.2, after some background on the penalized Nsa, we introduce a new algorithm called over-penalized NSa.

Section 3 is devoted to the main results: in Theorem 3.2, we establish an upper-bound of the *pseudo-regret* \bar{R}_n for the over-penalized algorithm in the two-armed case and also show a weaker result for the penalized NSa.

In this section, we also extend to the multi-armed case some existing convergence and rate of convergence results of the two-armed algorithm. In the “critical” case (see below for details), the normalized algorithm converges in distribution toward a PDMP (Piecewise Deterministic Markov Process). We develop a careful study of its ergodicity and bounds on the rate of convergence to equilibrium are established. It uses a non-trivial coupling strategy to derive explicit rates of convergence in Wasserstein and total variation distance. The dependence of these rates are made explicit with the several parameters of the initial Bandit problem.

The rest of the paper is devoted to the proofs of the main results: Section 4 is dedicated to the regret analysis, and Section 5 establishes the weak limit of the rescaled multi-armed bandit algorithm. Finally, Section 6 includes all the proofs of the ergodic rates.

2 Definitions of the NS algorithms

2.1 Crude NSa and regret

The crude NSa (1) is rather simple: it defines a $(\mathcal{F}_n)_{n \geq 0}$ Markov chain $(X_n)_{n \geq 0}$ and I_n is a random variable satisfying:

$$\mathbb{P}(I_{n+1} = 1 | \mathcal{F}_n) = X_n \quad \text{and} \quad \mathbb{P}(I_{n+1} = 2 | \mathcal{F}_n) = 1 - X_n$$

The arm I_{n+1} is selected at step $n+1$ with the current distribution $(X_n, 1 - X_n)$ and is evaluated. In the event of success, the weight of the arm I_{n+1} is increased and the weight of the other arm is decreased by the same quantity. The algorithm can be rewritten in a more concise form as:

$$X_{n+1} = X_n + \gamma_{n+1}(\mathbb{1}_{I_{n+1}=1} - X_n)A_{n+1}^{I_{n+1}}. \quad (7)$$

The arm i at step n succeeds with the probability $p_i = \mathbb{P}(A_n^i = 1)$ and we suppose *w.l.o.g.* that $p_1 > p_2$ so that the arm 1 is the optimal one.

As pointed in (6), we obtain that

$$\bar{R}_n = (p_1 - p_2) \mathbb{E} \left[\sum_{k=1}^n (1 - X_k) \right].$$

This formula is important regarding the fallibility of an algorithm. In particular, it is shown in [15] that for any choice $\gamma_n = C(n+1)^{-\alpha}$ with $\alpha \in (0, 1)$ and $C > 0$ or $\gamma_n = C/(n+1)$ with $C > 1$, the NSa (7) may be fallible: some parameters (p_1, p_2) exist such that $(X_n)_{n \geq 0}$ a.s. converges to a binary random variable X_∞ with $\mathbb{P}(X_\infty = 0) > 0$. In this situation, for large enough n , we have:

$$\bar{R}_n \gtrsim (p_1 - p_2) \mathbb{P}(X_\infty = 0) \times n \gg \sqrt{n}$$

It can easily be concluded that this method cannot induce a competitive policy since some “bad” values of the probabilities (p_1, p_2) generate a linear regret.

2.2 Penalized and over-penalized two-armed NSa

Penalized NSa. A major difference between the crude NSa and its penalized counterpart introduced in [16] relies on the exploitation of the failure of the selected arms. The crude NSa (1) only uses the sequence of successes to update the probability distribution $(X_n, 1 - X_n)$ since the value of X_n is modified *iff* $A_n^{I_n} = 1$. In contrast, the penalized NSa (2) also uses the information generated by a potential failure of the arm I_{n+1} . More precisely, in the event of success of the selected arm I_{n+1} , this penalized NSa mimics the crude NSa, whereas in the case of failure, the weight of the selected arm is now multiplied (and thus decreased) by a factor $(1 - \gamma_{n+1}\rho_{n+1})$ (whereas the probability of drawing the other arm is increased by the corresponding quantity). For the penalized NSa, the update formula of $(X_n)_{n \geq 1}$ can be written in the following way:

$$\begin{aligned} X_{n+1} = & X_n + \gamma_{n+1} [\mathbb{1}_{I_{n+1}=1} - X_n] A_{n+1}^{I_{n+1}} \\ & - \gamma_{n+1}\rho_{n+1} [X_n \mathbb{1}_{I_{n+1}=1} - (1 - X_n) \mathbb{1}_{I_{n+1}=2}] (1 - A_{n+1}^{I_{n+1}}). \end{aligned} \quad (8)$$

Over-penalized NSa. In view of the minimization of the regret, we will show that it may be useful to reinforce the penalization. For this purpose, we introduce a slightly “over-penalized” NSa where a player is also (slightly) penalized if it wins:

- If player 1 wins, then with probability $1 - \sigma$ it is penalized by a factor $\gamma_{n+1}\rho_{n+1}X_n$.
- If player 2 wins, then with probability $1 - \sigma$ arm 1 is increased by a factor of $\gamma_{n+1}\rho_{n+1}(1 - X_n)$.

The over-penalized-NSa can be written as follows

$$\begin{aligned} X_{n+1}^\sigma = & X_n^\sigma + \gamma_{n+1} [\mathbb{1}_{I_{n+1}=1} - X_n^\sigma] A_{n+1}^{I_{n+1}} \\ & - \gamma_{n+1}\rho_{n+1} [X_n^\sigma \mathbb{1}_{I_{n+1}=1} - (1 - X_n^\sigma) \mathbb{1}_{I_{n+1}=2}] (1 - A_{n+1}^{I_{n+1}} B_{n+1}^\sigma) \end{aligned} \quad (9)$$

where $(B_n^\sigma)_n$ is a sequence of i.i.d. r.v. with a Bernoulli distribution $\mathcal{B}(\sigma)$, meaning that $\mathbb{P}(B_n^\sigma = 0) = 1 - \sigma$. Moreover, these r.v. are independent of $(A_n^j)_{n,j}$ and in such a way that for all $n \in \mathbb{N}$, B_n^σ and I_n are also independent. It should be noted that

$$1 - A_n^{I_n} B_n^\sigma = [1 - A_n^{I_n}] + A_n^{I_n} (1 - B_n^\sigma).$$

In fact, this slight over-penalization of the successful arm (with probability σ) can be viewed as an additional statistical excitation which helps the stochastic algorithm to escape from local traps. The case $\sigma = 1$ corresponds to the penalized NSa (8), whereas when $\sigma = 0$, the arm is always penalized when it plays. In particular, this modification implies that the increment of X_n^σ is slightly weaker than in the previous case when the selected arm wins.

Asymptotic convergence of the penalized NSa. Before stating the main results, we need to understand which regret \bar{R}_n could be reached by penalized and over-penalized NSa. We recall (in a slightly less general form) the convergence results of Proposition 3, Theorems 3 and 4 of [16].

Theorem 2.1 (Lamberton & Pages, [16]). *Let $0 \leq p_2 < p_1 \leq 1$ and $\gamma_n = \gamma_1 n^{-\alpha}$ and $\rho_n = \rho_1 n^{-\beta}$ with $(\alpha, \beta) \in (0, +\infty)$ and $(\gamma_1, \rho_1) \in (0, 1)^2$. Let $(X_n)_n$ be the algorithm given by (8).*

i) *If $0 < \beta \leq \alpha$ and $\alpha + \beta \leq 1$, the penalized two-armed bandit is infallible.*

ii) *Furthermore, if $0 < \beta < \alpha$ and $\alpha + \beta < 1$, then $\frac{1 - X_n}{\rho_n} \rightarrow \frac{1 - p_1}{p_1 - p_2}$ a.s.*

iii) If $\alpha = \beta \leq 1/2$ and $g = \gamma_1/\rho_1$: $\frac{1 - X_n}{\rho_n} \xrightarrow{w^*} \mu$, where $\xrightarrow{w^*}$ stands for the convergence in distribution and μ is the stationary distribution of the PDMP whose generator \mathcal{L} acts on $\mathcal{C}_c^1(\mathbb{R}_+)$ as

$$\forall f \in \mathcal{C}_c^1(\mathbb{R}_+) \quad \mathcal{L}f(y) = p_2 y \frac{f(y+g) - f(y)}{g} + (1 - p_1 - p_1 y) f'(y).$$

In view of Theorem 2.1, we can use formula (6) to obtain

$$\bar{R}_n = (p_1 - p_2) \sum_{k=1}^n \rho_k \mathbb{E} \left(\frac{1 - X_k}{\rho_k} \right). \quad (10)$$

We then obtain the key observation

$$\sup_{n \in \mathbb{N}} \mathbb{E} \left[\frac{1 - X_n}{\rho_n} \right] \leq C < +\infty \implies \bar{R}_n \leq C(p_1 - p_2) \sum_{k=1}^n \rho_k, \quad (11)$$

where C is a constant that may depend on p_1 and p_2 . According to Theorem 2.1, it seems that the potential optimal choice corresponds to the one of (iii). Indeed, the infallibility occurs only when $\alpha \geq \beta$ and $\alpha + \beta \leq 1$ and Equation (10) suggests that β should be chosen as large as possible to minimize the r.h.s. of (11), leading to $\alpha = \beta = 1/2$. This is why in the following, we will focus on the case:

$$\gamma_n = \frac{\gamma_1}{\sqrt{n}} \quad \text{and} \quad \rho_n = \frac{\rho_1}{\sqrt{n}}. \quad (12)$$

2.3 Over-penalized multi-armed NSa

We generalize the definition of the penalized and over-penalized NSa to the d -armed case, with $d \geq 2$. Let $p = (p_1, \dots, p_d) \in (0, 1)^d$ and assume that $A_n^j \sim \mathcal{B}(p_j)$ (p_i the probability of success of arm i). The over-penalized NSa recursively defines a sequence of probability measures on $\{1, \dots, d\}$ denoted by $(\Pi_n)_{n \geq 1}$ where $\Pi_n = (X_n^1, \dots, X_n^d)$. At step n , the arm I_{n+1} is sampled according to the discrete distribution X_n and then tested through the computation of $A_{n+1}^{I_{n+1}}$. Setting $j \in \{1, \dots, d\}$, the multi-armed NSa is defined by:

$$\begin{aligned} X_{n+1}^j &= X_n^j + \gamma_{n+1} [\mathbb{1}_{I_{n+1}=j} - X_n^j] A_{n+1}^{I_{n+1}} \\ &\quad - \gamma_{n+1} \rho_{n+1} X_n^{I_{n+1}} (1 - A_{n+1}^{I_{n+1}} B_{n+1}^\sigma) \left[\mathbb{1}_{I_{n+1}=j} - \frac{1 - \mathbb{1}_{I_{n+1}=j}}{d-1} \right]. \end{aligned} \quad (13)$$

In contrast with the two-armed case, we have to choose how to distribute the penalty to the other arms when $d > 2$. The (natural) choice in (13) is to divide it fairly, *i.e.*, to spread it uniformly over the other arms. Note that alternative algorithms (not studied here) could be considered.

3 Main Results

3.1 Regret of the over-penalized two-armed bandit

First, we provide some uniform upper-bounds for the two-armed σ -over-penalized NSa. Our main result is Theorem 3.2. Before stating it, we choose to state a new result when $\sigma = 1$, *i.e.* for the “original” penalized NSa introduced in [16].

Theorem 3.1. *Let $(X_n)_{n \geq 0}$ be the two-armed penalized NSa defined by (8) with $(\gamma_n, \rho_n)_{n \geq 1}$ defined by (12) with $(\gamma_1, \rho_1) \in (0, 1)^2$. Then, for every $\delta \in (0, 1)$, a positive C_δ exists such that:*

$$\forall n \in \mathbb{N}, \quad \sup_{(p_1, p_2) \in [0, 1], p_2 \leq p_1 \wedge (1-\delta)} \bar{R}_n \leq C_\delta \sqrt{n}.$$

Remark 3.1. The upper bound of the original penalized-NS algorithm is not completely uniform. From a theoretical point of view, there is not enough penalty when p_2 is too large, which in turn generates a deficiency of the mean-reverting effect for the sequence $((1 - X_n)/\rho_n)_{n \geq 1}$ when X_n is close to 0. In other words, the trap of the stochastic algorithm near 0 is not enough repulsive and Figure 1 below shows that this problem also appears numerically and suggests a logarithmic explosion of $\sup_{p_1 < p_2} \bar{R}_n/\sqrt{n}$.

This explains the interest of the over-penalization, illustrated by the next result, which is the main theorem of the paper.

Theorem 3.2. Let $(X_n)_{n \geq 0}$ be the two-armed σ -over-penalized NSa defined by (9) with $\sigma \in [0, 1]$ and $(\gamma_n, \rho_n)_{n \geq 1}$ defined by (12) with $(\gamma_1, \rho_1) \in (0, 1)^2$. Then,

(a) A $C_\sigma(\gamma_1, \rho_1)$ exists such that:

$$\forall n \in \mathbb{N}, \quad \sup_{(p_1, p_2) \in [0, 1], p_2 < p_1} \bar{R}_n \leq C_\sigma(\gamma_1, \rho_1) \sqrt{n}.$$

(b) Furthermore, the choice $\sigma = 0$, $\gamma_n = 2.63\rho_n = 0.89/\sqrt{n}$ yields

$$\forall n \in \mathbb{N}, \quad \sup_{(p_1, p_2) \in [0, 1], p_2 < p_1} \bar{R}_n \leq 31.1\sqrt{2n}. \quad (14)$$

Remark 3.2. At the price of technicalities, C_σ could be made explicit in terms of γ_1 and ρ_1 for every $\sigma > 0$. The second bound is obtained by an optimization of $C_0(\gamma_1, \rho_1)$ (see (38) and below).

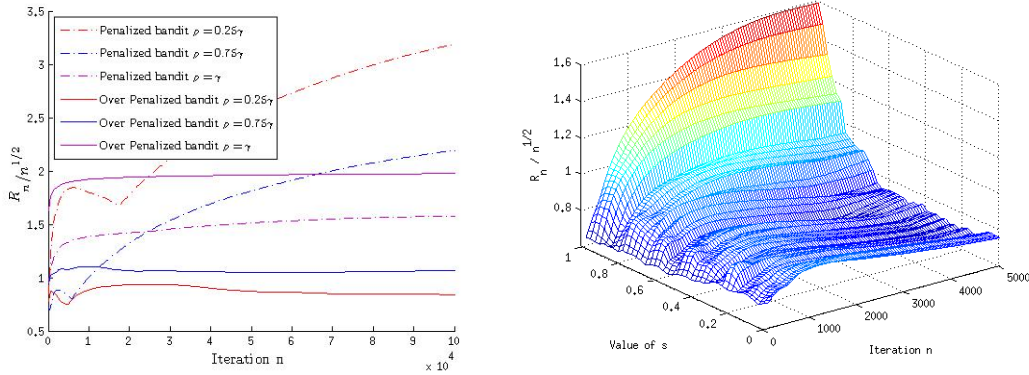


Figure 1: Evolution of $n \mapsto \sup_{(p_1, p_2) \in [0, 1], p_2 < p_1} \frac{\bar{R}_n}{\sqrt{n}}$ for the over-penalized algorithm (with $\sigma = 0$) and comparison with EXP3 and KL-UCB.

Figure 1 presents on the left side a numerical approximation of $n \mapsto \sup_{p_2 < p_1} R_n/\sqrt{n}$ for the penalized

and over-penalized algorithms. The continuous curves indicate that the upper bound $31.1\sqrt{2}$ in Theorem 3.2 is not sharp since the over-penalized NSa satisfies a uniform upper-bound on the order of $0.9\sqrt{n}$. This bound is obtained with a small σ (as pointed in Theorem 3.2), and $\gamma_n = \frac{1}{\sqrt{4+n}} = 4\rho_n$ (red line in Figure 1 (left)), suggesting that the rewards should *always* be over-penalized with $\rho_n = \frac{\gamma_n}{4}$.

The right-hand side of Figure 1 focuses on the behavior of the regret with σ . The map $(n, \sigma) \mapsto \sup_{p_1 < p_2} R_n/\sqrt{n}$ confirms the influence of the over-penalization and indicates that to obtain optimal performances for the cumulative regret, we should use a low value of σ between 0 and 3/5. The importance of this choice of σ seems relative since the behaviour of the over-penalized bandit is

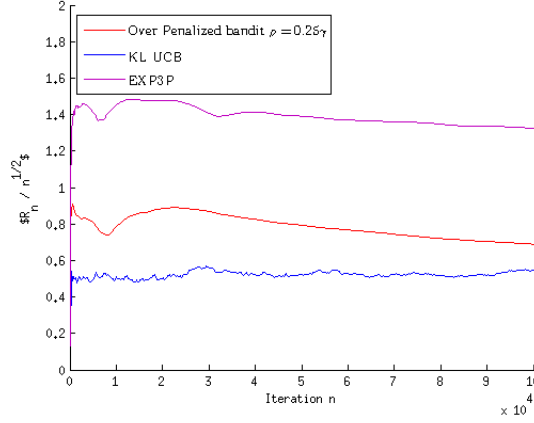


Figure 2: Evolution of $n \mapsto \sup_{(p_1, p_2) \in [0, 1], p_2 \leq p_1} \frac{\bar{R}_n}{\sqrt{n}}$ for the over-penalized algorithm (with $\sigma = \frac{1}{4}$) and comparison with EXP3 and KL-UCB.

stable on this interval. The best numerical choice is attained for $\sigma = 1/4$ and $\rho_n = \frac{1}{4}\gamma_n$ and permits to achieve a long-time behavior of \bar{R}_n/\sqrt{n} of the order $3/4$ (see Figure 2, red line).

Finally, the statistical performances of the over-penalized NSa are compared with some classical bandit algorithms: KL-UCB algorithm (see *e.g.* [9] and the references therein) and EXP3 (see [5]). These two algorithms are anytime policies that are known to be minimax optimal with a cumulative minimax regret of the order \sqrt{n} . Figure 2 shows that the performances of the over-penalized NSa are located between the one of the KL-UCB algorithm and of the EXP3 algorithm (our simulations suggest that the uniform bounds of KL-UCB and EXP3 are respectively $1/2$ and $3/2$). Also, it is worth noting that the simulation cost of the over-penalized NSa is strongly weaker than the initial UCB algorithm (the phenomenon is increased when compared to KL-UCB, which requires an additional difficulty for the computation of the upper confidence bound at each step): the same amount of Monte-Carlo simulations for the over-penalized NSa is almost hundred times faster than the KL-UCB runs in equivalent numerical conditions.

3.2 Convergence of the multi-armed over-penalized bandit

We first extend Theorem 2.1 of [16] to the over-penalized NSa in the multi-armed situation. The result describes the pointwise convergence.

Proposition 3.1 (Convergence of the multi-armed over-penalized bandit). *Consider $p_d \leq \dots \leq p_2 < p_1$ and $\gamma_n = \gamma_1 n^{-\alpha}$, $\rho_n = \rho_1 n^{-\beta}$ with $(\alpha, \beta) \in (0, +\infty)$ and $(\gamma_1, \rho_1) \in (0, 1)^2$. Algorithm (8) with $\sigma \in (0, 1]$ satisfies*

i) *If $0 < \beta \leq \alpha$ and $\alpha + \beta \leq 1$, then $\lim_{n \rightarrow +\infty} \Pi_n = (1, 0, \dots, 0)$ a.s.*

ii) *Furthermore, if $0 < \beta < \alpha$ and $\alpha + \beta < 1$, then:*

$$\forall i \in \{2, \dots, d\}, \quad \frac{X_n^i}{\rho_n} \longrightarrow \frac{1 - \sigma p_1}{(d-1)(p_1 - p_i)} \quad a.s.$$

Proposition 3.2 provides a description of the behavior of the *normalized* NSa while considering $Y_{n,j} = \frac{X_{n,j}}{\rho_n}$. It states that $(Y_{n,\cdot})_{n \geq 0}$ converges to the dynamics of a *Piecewise Deterministic Markov Process* (referred to as PDMP below).

Proposition 3.2 (Weak convergence of the over-penalized NSa). *Under the assumptions of Proposition 3.1, if $\alpha = \beta \leq 1/2$ and $g = \gamma_1/\rho_1$, then:*

$$\frac{1}{\rho_n} (X_{n,2}, \dots, X_{n,d}) \xrightarrow{w^*} \mu_d,$$

where μ_d is the (unique) stationary distribution of the Markov process whose generator \mathcal{L}_d acts on compactly supported functions f of $\mathcal{C}^1((\mathbb{R}_+)^{d-1})$ as follows:

$$\begin{aligned} \mathcal{L}_d f(y_2, \dots, y_d) &= \sum_{i=2, \dots, d} \frac{p_i y_i}{g} (f(y_2, \dots, y_i + g, \dots, y_d) - f(y_2, \dots, y_i, \dots, y_d)) \\ &+ \sum_{i=2, \dots, d} \left(\frac{1 - \sigma p_1}{d-1} - p_1 y_i \right) \partial_i f(y_2, \dots, y_d). \end{aligned} \quad (15)$$

3.3 Ergodicity of the limiting process

In this section, we focus on the long time behavior of the limiting Markov process that appears (after normalization) in Proposition 3.2. As mentioned before, this process is a PDMP and its long time behavior can be carefully studied with some arguments in the spirit of [6]. We also learned about the existence of a close study in the PhD thesis of Florian Bouguet (some details may be found in [8]). Such properties are stated for both the one-dimensional and the multidimensional cases.

3.3.1 One-dimensional case

Setting

$$a = 1 - p_1, \quad b = p_1, \quad g = \frac{\gamma_1}{\rho_1}, \quad c = \frac{p_2}{g},$$

the generator \mathcal{L} given by Proposition 3.2 may be written as:

$$\forall f \in \mathcal{C}^1(\mathbb{R}_+^*, \mathbb{R}) \quad \mathcal{L}f(x) = \underbrace{(a - bx)f'(x)}_{\text{deterministic part}} + \underbrace{cx}_{\text{jump rate}} \underbrace{(f(x+g) - f(x))}_{\text{jump size}}. \quad (16)$$

In what follows, we will assume that a, b, c and g are positive numbers. We can see in \mathcal{L} two parts. On the one hand, the deterministic flow that guides the PDMP between the jumps is given by:

$$\begin{cases} \partial_t \phi(x, t) &= (a - bx) \partial_x \phi(x, t) \\ \phi(x, 0) &= x \in \mathbb{R}_+^* \end{cases}$$

so that

$$\phi(x, t) = \frac{a}{b} + \left(x - \frac{a}{b} \right) e^{-bt}.$$

Hence, if $x > \frac{a}{b}$ (resp. $x < \frac{a}{b}$), $t \mapsto \phi(x, t)$ decreases (resp. increases) and converges exponentially fast to $\frac{a}{b}$.

On the other hand, the PDMP possesses some positive jumps that occur with a Poisson intensity “ $c \cdot x$ ”, whose size is deterministic and equals to g .

From the finiteness and positivity of g , it is easy to show that for every positive starting point, the process is *a.s.* well-defined on \mathbb{R}_+ , positive and does not explode in finite time. The fact that the size of the jumps is deterministic is less important and what follows could easily be generalized to a random size g (under adapted integrability assumptions). In Figure 3 below, some paths of the process are represented with different values of the parameters.

3.3.2 Convergence results

As pointed out in Figure 3, the long-time behavior of the process certainly depends on the relationship between the mean-reverting effect generated by “ $-bx$ ” and the frequency and size of the jumps.

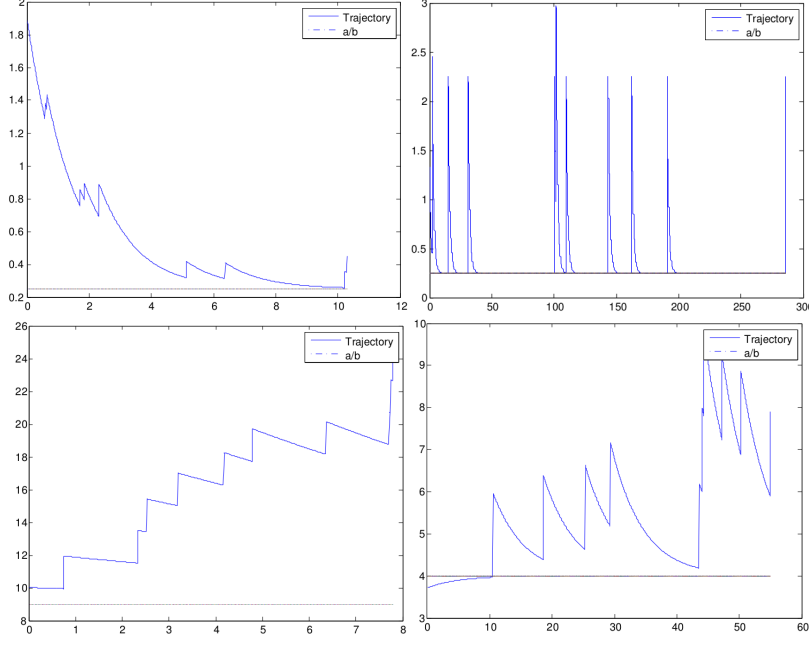


Figure 3: Exact simulation of trajectories of a process driven by (16) when $g = 0.1, a = 0.2, b = 0.8, c = 0.2$ (top left) $g = 2, a = 0.2, b = 0.8, c = 0.1$ (top right), $g = 2, a = 0.9, b = 0.9, c = 0.15$ (bottom left) and $g = 2, a = 0.8, b = 0.2, c = 0.05$ (bottom right).

Invariant measure The process (16) possesses a unique invariant distribution if $b - cg > 0$. Actually, the existence is ensured by the fact that $V(x) = x$ is a Lyapunov function for the process since

$$\forall x \in \mathbb{R}_+^*, \quad \mathcal{L}V(x) = a - (b - cg)x = a - (b - cg)V(x)$$

Among other arguments, the uniqueness is ensured by Theorem 3.3 (the convergence in Wasserstein distance of the process toward the invariant distribution implies in particular its uniqueness). We denote it by μ_∞ below. It could also be shown that $\text{Supp}(\mu_\infty) = (a/b, +\infty)$, that the process is strongly ergodic on $(a/b, +\infty)$ (see [12] for some background) and that if $b - cg > 0$, the process explodes when $t \rightarrow +\infty$ (this case corresponds to the bottom left-hand side of Figure 3). Finally, it should be noted that for the limiting PDMP of the bandit algorithm,

$$b - cg = p_1 - p_2 = \pi$$

and thus, the ergodicity condition coincides with the positivity of π .

Wasserstein results We aim to derive rates of convergence for the PDMP toward μ_∞ for two distances, namely the Wasserstein distance and the total variation distance. Rather different ways to obtain such results exist using coupling arguments or PDEs. We use coupling techniques here that are consistent with the work of [7] and [10]. Before stating our results, let us recall that the p -Wasserstein distance is defined for any probability measures μ and ν on \mathbb{R}^d by:

$$\mathcal{W}_p(\mu, \nu) = \inf \left\{ \mathbb{E}(|X - Y|^p)^{\frac{1}{p}} \mid \mathcal{L}(X) = \mu, \mathcal{L}(Y) = \nu \right\}.$$

Designating μ_0 as the initial distribution of the PDMP and μ_t as its law at time t , we now state the main result on the PDMP in dimension one driven by (16).

Theorem 3.3 (One dimensional PDMP). *Let $p \geq 1$ and denote for every $t \geq 0$ $\mu_t := \mathcal{L}(X_t^{\mu_0})$ where $(X_t^{\mu_0})$ is a Markov process driven by (16) with initial distribution μ_0 (with support included*

in \mathbb{R}_+^*). If $p = 1$, we have

$$\left| \int x(\mu_0 - \mu_\infty)(dx) \right| e^{-\pi t} \leq \mathcal{W}_1(\mu_t, \mu_\infty) \leq \mathcal{W}_1(\mu_0, \mu_\infty) e^{-t\pi}$$

and if $p > 1$, a constant γ_p exists such that

$$\mathcal{W}_p(\mu_t, \mu_\infty) \leq \gamma_p e^{-\frac{t\pi}{p}}.$$

where $(\gamma_p)_{p \geq 1}$ satisfies the recursion $\gamma_p^p = \gamma_{p-1}^{p-1}[pa + (1+g)^p]$.

Remark 3.3. If $p = 1$, the lower and upper bounds imply the optimality of the rate obtained in the exponential. For $p > 1$, the optimality of the exponent $e^{-\pi t/p}$ is still an open question.

We now give a corollary for the limiting process that appears in Proposition 3.2.

Corollary 3.1 (Multi-dimensional PDMP). *Let $(Y_t)_{t \geq 0}$ be the PDMP driven by (15) with initial distribution $\mu_0 \in (\mathbb{R}_+^*)^d$. Then, the conclusions of Theorem 3.3 hold with $\pi = p_1 - p_2$.*

The proof is almost obvious due to the “tensorized” form of the generator \mathcal{L}_d . Actually, for every starting point $y = (y_2, \dots, y_d)$, all the coordinates $(Y_t^i)_{t \geq 0}$ are independent one-dimensional PDMPs with generator \mathcal{L} defined by (16) with

$$a_i = \frac{1 - \sigma p_1}{d - 1}, \quad b_i = p_1 \quad \text{and} \quad c_i = p_i/g. \quad (17)$$

The result then easily follows from Theorem 3.3 with a global rate given by $\min\{b_i - c_i g, i = 2, \dots, d\} = p_1 - p_2$. The details are left to the reader.

3.4 Total variation results

When some bounds are available for the Wasserstein distance, a classical way to deduce an upper bound of the total variation is to build a two-step coupling. In the first step, a Wasserstein coupling is used to bring the paths sufficiently close (with a probability controlled by the Wasserstein bound). In a second step, we use a total variation coupling to try to stick the paths with a high probability. In our case, the jump size is deterministic and sticking the paths implies a non trivial coupling of the jump times. Some of the ideas to obtain the results below are in the spirit of [7], who follows this strategy for the TCP process.

Theorem 3.4. *Let μ_0 be a starting distribution with moments of any order. Then, for every $\varepsilon > 0$, a $C_\varepsilon > 0$ exists such that:*

$$\|\mu_0 P_t - \mu_\infty P_t\|_{TV} \leq C_\varepsilon e^{-(\alpha\pi - \varepsilon)t} \quad \text{with } \alpha = \frac{1}{2 + \frac{b\pi}{ac}}.$$

Once again, this result can be extended to the multi-armed case.

Corollary 3.2. *Let $(Y_t)_{t \geq 0}$ be the PDMP driven by (15) with initial distribution $\mu_0 \in (\mathbb{R}_+^*)^d$. Then, the conclusions of Theorem 3.4 hold with $\alpha\pi$ replaced by:*

$$\sum_{i=2}^d \frac{1}{2 + \frac{b_i \pi_i}{a_i c_i}} \pi_i$$

where $\pi_i = p_1 - p_i$ and a_i, b_i and c_i are defined by (17).

The proof of this result is based on the remark that follows Corollary 3.1. Owing to the “tensorization” property, the probability for coupling all the coordinates before time t is essentially the product of the probabilities of the coupling of each coordinate. Once again, the details of this corollary are left to the reader.

4 Proof of the regret bound (Theorems 3.1 and 3.2)

This section is devoted to the study of the regret of the penalized two-armed bandit procedure described in Section 2. We will mainly focus on the proof of the explicit bound given in Theorem 3.2(b) and we will give the main ideas for the proofs of Theorems 3.1 and 3.1(a).

4.1 Notations

In order to lighten the notations, X_n^1 will be summarized by X_n , so that $X_n^2 = 1 - X_n$. The proofs are then strongly based on a detailed study of the behavior of the (positive) sequence $(Y_n)_{n \geq 1}$ defined by

$$\forall n \geq 1 \quad Y_n = \frac{1 - X_n}{\gamma_n}. \quad (18)$$

As we said before, we will consider the following sequences $(\gamma_n)_{n \geq 1}$ and $(\rho_n)_{n \geq 1}$ below:

$$\forall n \geq 1, \quad \gamma_n = \frac{\gamma_1}{\sqrt{n}} \quad \text{and} \quad \rho_n = \frac{\rho_1}{\sqrt{n}} = \tilde{\rho}_1 \gamma_n \quad \text{and} \quad \tilde{\rho}_1 = \frac{\rho_1}{\gamma_1},$$

where γ_1 and ρ_1 are constants in $(0, 1)$ that will be specified later. In the meantime, we also define:

$$\pi = p_1 - p_2 \in (0, 1).$$

With this setting, the pseudo-regret is

$$\bar{R}_n = \pi \sum_{n=1}^n \gamma_n \mathbb{E}[Y_n].$$

It should be noted here that we have substituted the division by ρ_n in (11) by a normalization with γ_n . This will be easier to handle in the sequel. The main issue now is to obtain a convenient upper bound for $\mathbb{E}[Y_n]$. More precisely, note that:

$$\forall n_0 \in \mathbb{N} \quad \forall n \leq n_0 - 1, \quad \bar{R}_n \leq \pi n \leq \pi \sqrt{n_0 - 1} \sqrt{n},$$

and conversely for every $n \geq n_0$,

$$\begin{aligned} \frac{\bar{R}_n}{\sqrt{n}} &\leq \pi \sqrt{n_0 - 1} + \pi \sup_{n \geq n_0} \mathbb{E}[Y_n] \frac{1}{\sqrt{n}} \sum_{n=n_0}^n \frac{\gamma_1}{\sqrt{k}} \\ &\leq \pi \left(\sqrt{n_0 - 1} + 2\gamma_1 \sup_{n \geq n_0} \mathbb{E}[Y_n] \right). \end{aligned} \quad (19)$$

Thus it is enough to derive an upper bound of $\mathbb{E}[Y_n]$ after an iteration n_0 that can be on the order of $1/\pi^2$. In particular, the “suitable” choice of n_0 will strongly depend on the value of π .

4.2 Evolution of $(Y_n)_{n \geq 1}$

Recursive dynamics of $(Y_n)_{n \geq 1}$. In order to understand the mechanism and difficulties of the penalized procedure, let us first roughly describe the behavior of the sequences $(X_n)_{n \geq 1}$ and $(Y_n)_{n \geq 1}$. According to (9),

$$\begin{aligned} \mathbb{E}[X_{n+1} | \mathcal{F}_n] &= X_n + \gamma_{n+1} X_n (1 - X_n) [p_1 - p_2] \\ &\quad + \gamma_{n+1} \rho_{n+1} [(1 - X_n)^2 (1 - \sigma p_2) - X_n^2 (1 - \sigma p_1)]. \end{aligned}$$

It can be observed that the drift term may be split into two parts, where the main part is the usual drift of NSa described by h defined by:

$$\forall x \in [0, 1], \quad h(x) = [p_1 - p_2]x(1 - x). \quad (20)$$

The second term comes from the penalization procedure and depends on σ . We set

$$\kappa_\sigma(x) = (1 - \sigma p_2)(1 - x)^2 - (1 - \sigma p_1)x^2. \quad (21)$$

As a consequence, we can write the evolution of $(X_n)_{n \geq 0}$ as follows:

$$1 - X_{n+1} = 1 - X_n - \gamma_{n+1} [h(X_n) + \rho_{n+1} \kappa_\sigma(X_n) + \Delta M_{n+1}], \quad (22)$$

where ΔM_{n+1} is a martingale increment. On the basis of the equation above, we easily derive that

$$\forall n \geq 1, \quad Y_{n+1} = Y_n (1 + \gamma_n (\epsilon_n - \pi X_n)) - \rho_{n+1} \kappa(X_n) + \Delta M_{n+1}$$

where

$$\epsilon_n = \frac{1}{\gamma_{n+1}} - \frac{1}{\gamma_n} = \frac{1}{\gamma_1} (\sqrt{n+1} - \sqrt{n}) \leq \frac{1}{2\gamma_1 \sqrt{n}} = \frac{\gamma_n}{2\gamma_1^2}. \quad (23)$$

It follows that the increments of $(Y_n)_{n \geq 1}$ are given by:

$$\Delta Y_{n+1} := Y_{n+1} - Y_n = \gamma_n \varphi_n(Y_n) - \Delta M_{n+1}$$

where the drift function φ_n acting on the sequence $(Y_n)_{n \geq 1}$ is defined as

$$\varphi_n(y) = \underbrace{y \times [\epsilon_n + \pi(\gamma_n y - 1)]}_{:= \varphi_n^1(y)} + \underbrace{\left(-\frac{\rho_{n+1}}{\gamma_n} \kappa_\sigma(1 - \gamma_n y) \right)}_{:= \varphi_n^2(y)}.$$

To better understand the underlying effects of the dynamical system, it should be recalled that the definition of the sequence $(Y_n)_{n \geq 1}$ implies that $Y_n \in [0, \gamma_n^{-1}]$ with $\gamma_n^{-1} \sim Cn^{1/2}$. Since we aim to obtain a uniform bound (over n) of $\mathbb{E}[Y_n]$, it is thus important to understand the behavior of the drift φ_n over $[0, \gamma_n^{-1}]$. In particular, it is of primary interest to see where the function φ_n is negative.

Crude NSa. When dealing with the crude bandit algorithm (*i.e.*, when $\rho_1 = 0$, see (7)), the drift is reduced to φ_n^1 . One can check that $\varphi_n^1(y)$ is negative *iff*

$$\epsilon_n - \pi(1 - \gamma_n y) < 0 \iff y \leq \gamma_n^{-1} - \frac{\epsilon_n}{\pi \gamma_n} \iff x \geq \frac{\epsilon_n}{\pi}$$

where $x = 1 - \gamma_n y$. This means that when x is close to 0 (in some sense depending on n , π and γ_1), φ_n^1 becomes positive and Y_n has a tendency to increase. In others words, the dynamical system $(Y_n)_{n \geq 1}$ has no *mean-reverting* when Y_n is far from 0. The fact that the crude bandit algorithm does not always converge to the good target can be understood as a consequence of this remark.

Penalized and Over-Penalized NSa. When the drift φ_n contains a non zero penalty, the second term $-\varphi_n^2$ may help the dynamics to not be repulsive when x is close to 0, *i.e.* when y is larger than $1/\gamma_n$. It can be checked that $\kappa_\sigma(0) = 1 - \sigma p_2$ and:

$$\lim_{n \rightarrow +\infty} \varphi_n(\gamma_n^{-1}) = \frac{1}{2\gamma_1^2} - \frac{\gamma_1}{\rho_1} (1 - \sigma p_2).$$

This quantity is negative under the condition:

$$1 - \sigma p_2 > \frac{\rho_1}{2\gamma_1^3}. \quad (24)$$

But, in order to obtain a uniform bound on the regret, this constraint must be satisfied independently of p_2 . When $\sigma = 1$, *i.e.* in the standardly penalized case, one remarks that for any choice of ρ_1 and γ_1 , this is only possible if $\rho_1/(2\gamma_1^3) > 1 - p_2$. At this stage, one can thus understand the

over-penalization as a way of controlling uniformly (in p_2) the negativity of φ_n far from $y = 0$ (see Figure 4).

In view of the main results, there are still two problems. The first one is that even in the over-penalized case, Inequality (24) implies some constraints on γ_1 and ρ_1 , which do not appear in Theorem 3.2. The second one which is more embarrassing for the study of $(\mathbb{E}[Y_n])_{n \geq 1}$ is that, near $y = 0$, φ_n is positive since $\varphi_n(0) = 1 - \sigma p_1$ ((see Figure 4)). This repulsive behavior near $y = 0$ can be understood as the counterpart induced by the penalization. In order to bypass the two previous problems, the main argument will be the *increase of exponent* (see next section) where we show that we can replace the study of $(\mathbb{E}[Y_n])_{n \geq 1}$ by the one of a sequence which both has a nicer behavior near $y = 0$ and alleviates the constraint (24).

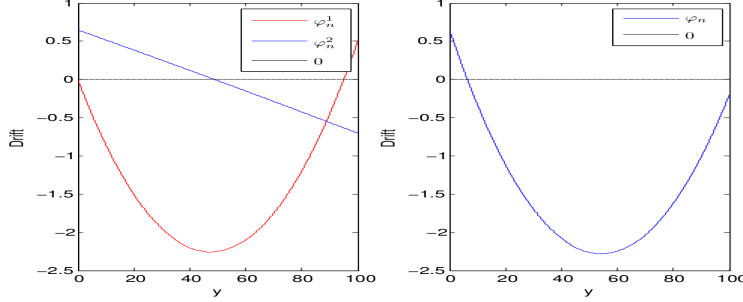


Figure 4: Drift decomposition (left) and global (right) when $y \in [0, \frac{1}{\gamma_n}]$ with $\gamma_1 = \rho_1 = 1$, $p_1 = 0.7$, $p_2 = 0.6$, $\sigma = 0.5$.

4.3 Increase of exponent

We introduce the sequence $(Z_n^{(r)})_{n \geq 0}$ defined by:

$$\forall n \geq 1 \quad Z_n^{(r)} = \frac{(1 - X_n)^r}{\gamma_n}. \quad (25)$$

At this stage, one can first remark that *a.s.*, for every $r \geq 1$, $Z_n^{(r)} \leq Z_n^{(1)} = Y_n$. One can thus guess that the difficulties tackled at the end of the previous section will be easier to overcome for $(\mathbb{E}[Z_n^{(r)}])_{n \geq 1}$ with $r > 1$. Of course, this remark has an interest if conversely, one is able to relate the control of $\mathbb{E}[Y_n]$ to those of $\mathbb{E}[Z_n^{(r)}]$, $r \geq 1$.

This is the purpose of Proposition 4.1 where taking advantage of the structure of the algorithm, one shows that for every $r \geq 1$, $\mathbb{E}[Z_n^{(r)}]$ can be controlled by a function of $\mathbb{E}[Z_n^{(r+1)}]$.

Let us define the bounded function h_r on $[0, 1]$:

$$\forall \gamma \in [0, 1] \quad h_r(\gamma) = \frac{(1 + \gamma)^r - 1 - r\gamma}{r\gamma^2}. \quad (26)$$

Proposition 4.1. *Let $r \in \mathbb{N}^*$, $\gamma_1 \in (0, 1)$ and $0 < \epsilon \leq \epsilon_0 = \frac{1}{3}$, and set*

$$n_0(\epsilon, \pi, \gamma_1) := \left\lceil \frac{1}{4\epsilon^2\gamma_1^2\pi^2} \right\rceil + 1. \quad (27)$$

Then, if $2\epsilon\gamma_1^2(r - \epsilon) \leq 1$,

$$\sup_{n \geq n_0} \mathbb{E}Z_n^{(r)} \leq \mathbb{E}Z_{n_0}^{(r)} + \frac{r}{\pi(r - \epsilon)} \left[\tilde{\rho}_1 + h_r(\gamma_{n_0}) + \pi \sup_{n \geq n_0} \mathbb{E}[Z_n^{(r+1)}] \right].$$

In particular, for $r = 1, 2$, the previous inequality holds for every $\gamma_1 \in (0, 1)$ and $\epsilon \in (0, 1/3]$.

Remark 4.1. Note that the above result induces some constraints on ϵ and γ . These constraints, which allow us to manage the constants of the inequality, are mainly adapted to the proof of Theorem 3.2 (b). In fact, in the proofs of Theorems 3.1 and 3.2 (a), we will need to rewrite the above property in a slightly different way (see Section 4.5 for details).

Proof For any integer $r > 0$ and $n \geq 0$, the binomial formula applied to (22) leads to

$$\begin{aligned} (1 - X_{n+1})^r &= (1 - X_n - \Delta X_{n+1})^r \\ &= (1 - X_n)^r - r(1 - X_n)^{r-1} \Delta X_{n+1} \\ &\quad + \sum_{j=0}^{r-2} \binom{r}{j} (1 - X_n)^j (-\Delta X_{n+1})^{r-j}. \end{aligned}$$

where $\sum_{\emptyset} = 0$ and $\Delta X_{n+1} = X_{n+1} - X_n = \gamma_{n+1}[h(X_n) + \rho_{n+1}\kappa_{\sigma}(X_n) + \Delta M_{n+1}]$. From the definition of h given in (20), we get

$$(1 - x)^{r-1}[h(x) + \rho_{n+1}\kappa_{\sigma}(x)] = \pi x(1 - x)^r + \rho_{n+1}\kappa_{\sigma}(x)(1 - x)^{r-1}.$$

If we define now

$$\beta_n^{(r)} = -r\rho_{n+1}(1 - X_n)^{r-1}\kappa_{\sigma}(X_n) + \frac{1}{\gamma_{n+1}} \sum_{j=0}^{r-2} \binom{r}{j} (1 - X_n)^j (-\Delta X_{n+1})^{r-j}, \quad (28)$$

we can then conclude using (25) that

$$\begin{aligned} Z_{n+1}^{(r)} &= Z_n^{(r)} \frac{\gamma_n}{\gamma_{n+1}} - \gamma_n r \pi X_n Z_n^{(r)} + \beta_n^{(r)} - r(1 - X_n)^{r-1} \Delta M_{n+1} \\ &= Z_n^{(r)} \left(1 + \gamma_n \left[\frac{1}{\gamma_{n+1}} - \frac{1}{\gamma_n} - r\pi X_n \right] \right) + \beta_n^{(r)} - r(1 - X_n)^{r-1} \Delta M_{n+1} \\ &= Z_n^{(r)} (1 + \gamma_n [\epsilon_n - r\pi X_n]) + \beta_n^{(r)} - r(1 - X_n)^{r-1} \Delta M_{n+1} \\ &= Z_n^{(r)} (1 + \gamma_n [\epsilon_n - r\pi]) + r\pi \gamma_n (1 - X_n) Z_n^{(r)} + \beta_n^{(r)} - r(1 - X_n)^{r-1} \Delta M_{n+1} \\ &= Z_n^{(r)} (1 + \gamma_n [\epsilon_n - r\pi]) + r\pi \gamma_n Z_n^{(r+1)} + \beta_n^{(r)} - r(1 - X_n)^{r-1} \Delta M_{n+1}. \end{aligned} \quad (29)$$

The formulation above is important: it exhibits a contraction of $(1 + \gamma_n [\epsilon_n - r\pi])$ on $Z_n^{(r)}$ that can be used jointly with an upper bound of $Z_n^{(r+1)}$ and a simple majorization of $\beta_n^{(r)}$. In this view, we study (28): $|\Delta X_{n+1}| \leq \gamma_{n+1}$ a.s. and (21) yields $|\kappa_{\sigma}(x)| \leq (1 - \sigma p_2)$. Now, with h_r given in (26), we get

$$\beta_n^{(r)} \leq r\tilde{\rho}_1 \gamma_{n+1} + \sum_{j=0}^{r-2} \binom{r}{j} (\gamma_{n+1})^{r-j-1} \leq r(\tilde{\rho}_1 + h_r(\gamma_{n+1})) \gamma_{n+1}.$$

For any $\epsilon \in (0, 1)$, we can see in (29) that the contraction coefficient can be useful as soon as n is large enough. More precisely, using (23), we see that

$$\epsilon_n \leq \epsilon \iff n \geq n_0(\epsilon, \pi, \gamma_1) := \left\lfloor \frac{1}{4\epsilon^2 \gamma_1^2 \pi^2} \right\rfloor + 1.$$

Then, for every $n \geq n_0(\epsilon, \pi, \gamma_1)$,

$$1 + \gamma_n [\epsilon_n - r\pi] \leq 1 - \alpha_r \gamma_n \quad \text{with} \quad \alpha_r = \pi(r - \epsilon).$$

In the sequel, we will omit the dependence of n_0 in $(\epsilon, \pi, \gamma_1)$ and will just use the notation n_0 . Also remark that under the condition $2\epsilon\gamma_1^2(r - \epsilon) \leq 1$, we have $\alpha_r \gamma_j < 1$ for every $\pi \in (0, 1)$ and

for every $j \geq n_0$ (one can in particular check that $2\epsilon\gamma_1^2(r-\epsilon) \leq 1$ is true for every $\epsilon \in (0, 1/3)$ and $\gamma_1 \in (0, 1)$ if $r = 1, 2$). Thus, by a simple recursion based on (29), one obtains for every $n \geq n_0 + 1$,

$$\mathbb{E}(Z_n^{(r)}) \leq \mathbb{E}(Z_{n_0}^{(r)}) \prod_{j=n_0}^{n-1} (1 - \alpha_r \gamma_j) + \sum_{j=n_0}^{n-1} \left(r\pi\gamma_j \mathbb{E}(Z_j^{(r+1)}) + \beta_j^{(r)} \right) \prod_{l=j}^{n-1} (1 - \alpha_r \gamma_l)$$

If we call $\Theta_r = r \left(\pi \sup_{j \geq n_0} \left(\mathbb{E}(Z_j^{(r+1)}) + \tilde{\rho}_1 + h_r(\gamma_j) \right) \right)$, an iteration of the previous inequality yields:

$$\mathbb{E}(Z_n^{(r)}) \leq \mathbb{E}(Z_{n_0}^{(r)}) + \Theta_r \sum_{j=n_0}^{n-1} \gamma_j \prod_{l=j}^{n-1} (1 - \alpha_r \gamma_l).$$

We aim to apply Lemma A.1 (deferred to the appendix section) to the last term. It is possible as soon as

$$n_0 \geq \frac{1}{(\alpha_r \gamma_1)^2}.$$

This last condition is fulfilled for any $r \geq 1$ when $\frac{1}{4\epsilon^2\gamma_1^2\pi^2} \geq \frac{1}{(1-\epsilon)^2\pi^2\gamma_1^2}$, i.e. when $\epsilon \leq 1/3$.

Then, by Lemma A.1, one deduces that $\forall \epsilon \leq 1/3$ and $\forall n \geq n_0$:

$$\sup_{n \geq n_0} \mathbb{E}Z_n^{(r)} \leq \mathbb{E}Z_{n_0}^{(r)} + \frac{r}{\pi(r-\epsilon)} \left[\tilde{\rho}_1 + h_r(\gamma_{n_0}) + \pi \sup_{n \geq n_0} Z_n^{(r+1)} \right].$$

□

On the basis of the last proposition and a recursive argument, we can now deduce the following key observations.

Corollary 4.1. *Assume that $\epsilon \in (0, 1/3)$, $\gamma_1 \in (0, 1)$ and that n_0 is defined in (27). Then,*

$$\begin{aligned} \sup_{n \geq n_0} \mathbb{E}[Y_n] &\leq \mathbb{E}[Z_{n_0}^{(1)}] + \frac{\mathbb{E}[Z_{n_0}^{(2)}]}{1-\epsilon} + \frac{1}{\pi(1-\epsilon)} \left[\tilde{\rho}_1 + \frac{\tilde{\rho}_1}{1-\epsilon/2} + \frac{1}{2(1-\epsilon/2)} \right] \\ &\quad + \frac{\sup_{n \geq n_0} \mathbb{E}Z_n^{(3)}}{(1-\epsilon)(1-\epsilon/2)}. \end{aligned} \quad (30)$$

Remark 4.2. *As in Proposition 4.1, this property is mainly written in view of Theorem 3.2 (b) where we only need to use the increase of exponent for $r = 1, 2$. For Theorems 3.1 and 3.2 (a) with $\sigma \in (0, 1)$, we will need to use it for large values of r .*

4.4 Bound for $(\mathbb{E}(Z_n^{(3)}))_{n \geq n_0}$

As seen in Corollary 4.1, our next task is to bound $\mathbb{E}(Z_n^{(3)})$ for $n \geq n_0$ to obtain a tractable application of Equation (30). Such a bound is reached through careful inspection of the increments $\Delta Z_{n+1}^{(3)} := Z_{n+1}^{(3)} - Z_n^{(3)}$.

Lemma 4.1 (Decomposition of $Z_n^{(3)}$). *For every $n \geq 1$,*

$$\mathbb{E}[\Delta Z_{n+1}^{(3)} | \mathcal{F}_n] = \gamma_{n+1}(1 - X_n)P_n(X_n) + \Delta R_n,$$

where for every $n \in \mathbb{N}$, P_n is a polynomial function defined by

$$\begin{aligned} P_n(x) &= \frac{(1-x)^2}{\gamma_{n+1}} (\epsilon_n - 3\pi x) - 3\tilde{\rho}_1(1-x)\kappa_\sigma(x) + 3(x(1-x)^2 p_1 + x^2(1-x)p_2) \\ &\quad + \gamma_{n+1}(-x(1-x)^2 p_1 + x^3 p_2), \end{aligned} \quad (31)$$

and if γ_1 and n_0 satisfy the assumptions of Proposition 4.1, then

$$\forall n \geq n_0, \quad \Delta R_n \leq (1 - \sigma p_2) [3\gamma_{n+1}\rho_{n+1}^2 + \gamma_{n+1}^2\rho_{n+1}^3].$$

Remark 4.3. • The keypoint is that $\gamma_k = \gamma_1 k^{-1/2}$ and, therefore, the series $\sum_{n \geq 1} \Delta R_k$ is uniformly bounded, regardless of the value of π . This will be enough to obtain a competitive upper bound of the regret. With the choice of n_0 given in (27), careful inspection of Lemma 4.1 leads to:

$$\sum_{k \geq n_0} \Delta R_k \leq 12\gamma_1^4 \rho_1^2 \epsilon \pi + \frac{16}{3} \gamma_1^8 \rho_1^3 \epsilon^3 \pi^3. \quad (32)$$

- As in Remark 4.2, it should be noted that for Theorems 3.1 and 3.2 (a) with $\sigma \in (0, 1)$, we will need to use such a development with some larger values of r (see the end of this section for details).

Proof. We again use Equation (29) and deduce that:

$$Z_{n+1}^{(3)} - Z_n^{(3)} = (1 - X_n)^3 (\epsilon_n - 3\pi X_n) - 3\tilde{\rho}_1 \gamma_{n+1} (1 - X_n)^2 \kappa_\sigma(X_n) \quad (33)$$

$$+ \frac{1}{\gamma_{n+1}} \sum_{j=0}^1 \binom{3}{j} (1 - X_n)^j (-\Delta X_{n+1})^{3-j} - 3(1 - X_n)^2 \Delta M_{n+1} \quad (34)$$

First, note that terms in Equation (33) are associated with the first two terms in the definition of P_n introduced in (31) up to a multiplication by $(1 - X_n)\gamma_{n+1}$.

Second, we can easily compute the expectations involved in the sum of Equation (34) since the events are all disjointed. On the one hand, when $j = 1$ we have

$$\begin{aligned} \frac{1}{\gamma_{n+1}} \mathbb{E}[(-\Delta X_{n+1})^2 | \mathcal{F}_n] &= \gamma_{n+1} \sigma (p_1 X_n (1 - X_n)^2 + p_2 (1 - X_n) X_n^2) \\ &+ \gamma_{n+1} (1 - \sigma) (p_1 X_n (1 - X_n - \rho_{n+1} X_n)^2 + p_2 (1 - X_n) (X_n - \rho_{n+1} (1 - X_n))^2) \\ &+ \gamma_{n+1} \rho_{n+1}^2 [X_n^3 (1 - p_1) + (1 - X_n)^3 (1 - p_2)]. \end{aligned}$$

Further computations yield:

$$\begin{aligned} \frac{1}{\gamma_{n+1}} \mathbb{E}[(-\Delta X_{n+1})^2 | \mathcal{F}_n] &= \gamma_{n+1} X_n (\sigma p_1 (1 - X_n)^2 + \sigma p_2 X_n^2) + \Delta A_n^{(1)} \\ &+ \underbrace{\gamma_{n+1} \rho_{n+1}^2 [X_n^3 (1 - \sigma p_1) + (1 - X_n)^3 (1 - \sigma p_2)]}_{:= \Delta R_n^{(1)}} \end{aligned}$$

with $\Delta A_n^{(1)} = -2\rho_{n+1}\gamma_{n+1}X_n(1 - X_n)(1 - \sigma)(X_n p_1 + (1 - X_n)p_2)$. On the other hand, we can also compute the term when $j = 0$:

$$\begin{aligned} \frac{1}{\gamma_{n+1}} \mathbb{E}[(-\Delta X_{n+1})^3 | \mathcal{F}_n] &= \gamma_{n+1}^2 X_n (1 - X_n) (p_2 X_n^2 - p_1 (1 - X_n)^2) \\ &+ \Delta A_n^{(2)} + \underbrace{\gamma_{n+1}^2 \rho_{n+1}^3 [X_n^4 (1 - \sigma p_1) - (1 - X_n)^4 (1 - \sigma p_2)]}_{:= \Delta R_n^{(2)}} \end{aligned}$$

with $\Delta A_n^{(2)} \leq 3\gamma_{n+1}^2 \rho_{n+1} (1 - \sigma) X_n (1 - X_n)^2 (\pi X_n + \rho_{n+1} (1 - X_n) p_2)$. Set $\Delta R_n^{(3)} = (1 - X_n) \Delta A_n^{(1)} + \Delta A_n^{(2)}$ and $\Delta R_n := 3(1 - X_n) \Delta R_n^{(1)} + \Delta R_n^{(2)}$. Plugging the previous controls into (34) yields

$$\mathbb{E}[\Delta Z_{n+1}^{(3)} | \mathcal{F}_n] \leq \gamma_{n+1} (1 - X_n) P_n(X_n) + \Delta R_n. \quad (35)$$

Note that $\Delta R_n^{(1)}$ can be upper bounded as follows:

$$3(1 - X_n) \Delta R_n^{(1)} \leq 3\gamma_{n+1} \rho_{n+1}^2 (1 - \sigma p_2) \max_{0 \leq t \leq 1} \left[\frac{1 - \sigma p_1}{1 - \sigma p_2} t^3 (1 - t) + (1 - t)^4 \right].$$

Since $1 - \sigma p_1 \leq 1 - \sigma p_2$, a study of the function shows that $at^3(1 - t) + (1 - t)^4$ when $a \in (0, 1)$ reaches its maximal value for $t = 0$. This leads to:

$$3(1 - X_n) \Delta R_n^{(1)} \leq 3\gamma_{n+1} \rho_{n+1}^2 (1 - \sigma p_2).$$

For $\Delta R_n^{(2)}$, we have $\Delta R_n^{(2)} \leq \gamma_{n+1}^2 \rho_{n+1}^3 (1 - \sigma p_2) \max_{0 \leq t \leq 1} \left[\frac{1 - \sigma p_1}{1 - \sigma p_2} t^4 - (1 - t)^4 \right]$, which involves an increasing function of t . Thus, we have

$$\Delta R_n^{(2)} \leq \gamma_{n+1}^2 \rho_{n+1}^3 (1 - \sigma p_1) \leq \gamma_{n+1}^2 \rho_{n+1}^3 (1 - \sigma p_2).$$

Finally, if γ_1 and n_0 satisfy the assumptions of Proposition 4.1, then for every $n \geq n_0$, $\gamma_n \leq 2/3$ and it follows that $\Delta R_n^{(3)} \leq 0$. The result follows according to Equation (35). \square

In order to bound $\sup_{n \geq n_0} \mathbb{E}(Z_n^{(3)})$, we now have to precisely study the polynomial function P_n and exhibit a mean reverting effect on its dynamics.

Proposition 4.2. *Let $\epsilon \in (0, \frac{1}{3})$, $\tilde{\rho}_1 \leq \frac{227}{232}$ and $\frac{1}{3\sqrt{2}(1-\sigma)\tilde{\rho}_1} \leq \gamma_1^2 \leq \frac{3}{2(1+\tilde{\rho}_1)}$. Then*

i) *The polynomial P_n given by (31) is negative on $[0, 1 - \frac{2(1+\tilde{\rho}_1)}{\pi} \gamma_{n+1}]$.*

ii) *$Z_n^{(3)}$ satisfies*

$$\begin{aligned} \sup_{n \geq n_0} \mathbb{E} Z_n^{(3)} &\leq \mathbb{E} Z_{n_0}^{(3)} + 12\gamma_1^4 \tilde{\rho}_1^2 \epsilon \pi + \frac{16}{3} \gamma_1^8 \tilde{\rho}_1^3 \epsilon^3 \pi^3 \\ &\quad + \frac{8\gamma_1^4 \epsilon (1 + \tilde{\rho}_1) [1 + (1 + \tilde{\rho}_1)[2 + 6\tilde{\rho}_1 + 12\gamma_1^2 \epsilon]]}{\pi}. \end{aligned}$$

Remark 4.4. *The above result is given under some technical conditions that will lead to a sharp explicit bound. Nevertheless, the reader has to keep in mind that in view of the condition on σ , the “universal” bound on $(\mathbb{E}(Z_n^{(3)}))_{n \geq n_0}$ is only accessible when $\sigma < 1$, i.e. in the over-penalized case. When $\sigma = 1$, some bounds will be attainable only if p_2 is not too large (see (24) for a similar statement when $r = 1$), and in order to alleviate the constraint on p_2 , it will be necessary to take a larger exponent than $r = 3$ (see Subsection 4.5 for details).*

Proof. We first provide the proof of i). The function P_n introduced in (31) is a third degree polynomial and for $n \geq n_0$:

$$\begin{aligned} P_n(0) &= \frac{\epsilon_n}{\gamma_{n+1}} - 3\tilde{\rho}_1 \kappa_\sigma(0) \\ &\leq \frac{\gamma_n}{2\gamma_1^2 \gamma_{n+1}} - 3\tilde{\rho}_1 (1 - \sigma p_2) \\ &\leq \frac{\sqrt{1 + n_0^{-1}}}{2\gamma_1^2} - 3\tilde{\rho}_1 (1 - \sigma p_2) \end{aligned}$$

Since $p_2 < 1$, this last quantity is negative if:

$$\rho_1 \gamma_1 \geq \frac{1}{3\sqrt{2}(1-\sigma)}. \quad (36)$$

In a same way, we can check that $P_n(1) = \gamma_{n+1} p_2 > 0$ and, therefore, P_n has one root in the interval $(0, 1)$. Careful inspection of the leading coefficient (designated $a_n x^3$) of P_n in (31) shows that:

$$a_n = \left[3(1 + \sigma \tilde{\rho}_1) - \frac{3}{\gamma_{n+1}} - \gamma_{n+1} \right] \pi.$$

The leading coefficient a_n is negative as soon as $3(1 + \sigma \tilde{\rho}_1) \leq \frac{3}{\gamma_{n+1}}$. Again, the choice of n_0 in (27) shows that this last condition is fulfilled as soon as

$$\frac{1}{\epsilon} \geq 2\gamma_1 \pi (\gamma_1 + \sigma \rho_1). \quad (37)$$

It should however be noted that we have assumed $\epsilon \in (0, 1/3]$ so that $\frac{1}{\epsilon} \geq 3$. As a consequence, (36) and (37) are satisfied as soon as $(\gamma_1, \tilde{\rho}_1)$ satisfies

$$\frac{1}{3\sqrt{2}(1-\sigma)\tilde{\rho}_1} \leq \gamma_1^2 \leq \frac{3}{2(1+\tilde{\rho}_1)}$$

Hence, if (36) and (37) hold, P_n possesses one root in $(-\infty, 0)$ and another one in $(1, +\infty)$. Consequently, P_n has a unique root in $(0, 1)$. We now consider:

$$\xi_n = \frac{2(1+\tilde{\rho}_1)}{\pi} \gamma_{n+1} := \xi \gamma_{n+1}.$$

We compute that:

$$\begin{aligned} & P_n(1 - \xi_n) \\ &= \frac{\xi_n^2}{\gamma_{n+1}} (\epsilon_n - 3\pi(1 - \xi_n)) - 3\tilde{\rho}_1 \xi_n [(1 - \sigma p_2) \xi_n^2 - (1 - \sigma p_1)(1 - \xi_n)^2] \\ & \quad + 3 [(1 - \xi_n) \xi_n^2 p_1 + (1 - \xi_n)^2 \xi_n p_2] + \gamma_{n+1} [(1 - \xi_n)^3 p_2 - \xi_n^2 (1 - \xi_n)]. \end{aligned}$$

Hence, replacing ξ_n by $\xi \gamma_{n+1}$ and simplifying by γ_{n+1} , we see that $P_n(1 - \xi_n)$ is negative when

$$\begin{aligned} & \overbrace{\frac{\xi^2 \epsilon_n}{(1 - \xi_n)} + 3\tilde{\rho}_1(1 - \sigma p_1)(1 - \xi_n)\xi + 3p_1\gamma_{n+1}\xi^2 + 3p_2(1 - \xi_n)\xi + p_2(1 - \xi_n)^2}^{:=A_n(\xi)} \\ & \leq \underbrace{3\pi\xi^2 + \frac{3\tilde{\rho}_1\gamma_{n+1}^2\xi^3(1 - \sigma p_2)}{1 - \xi_n}}_{:=B_n(\xi)} + \gamma_{n+1}^2\xi^2. \end{aligned}$$

From (23), we know that $\epsilon_n \leq \frac{\gamma_{n+1}}{2\gamma_1^2}$, and $1 - \xi_n \leq 1$ thus

$$A_n(\xi) \leq \xi^2 \gamma_{n+1} \left(\frac{1}{2\gamma_1^2(1 - \xi_n)} + 3p_1 \right) + 3\xi(\tilde{\rho}_1 + 1) + 1$$

In the meantime, we will use the simple lower bound $B_n(\xi) \geq 3\pi\xi^2$. We can check that $1 - \xi_n = 1 - \frac{2(1+\tilde{\rho}_1)\gamma_{n+1}}{\pi} \geq 1 - 4\epsilon(1 + \tilde{\rho}_1)\gamma_1^2$ since $\gamma_{n_0} \leq 2\epsilon\gamma_1^2\pi$. Thus

$$\begin{aligned} & A_n \left(\frac{2(1+\tilde{\rho}_1)}{\pi} \right) \\ & \leq \frac{4(1+\tilde{\rho}_1)^2}{\pi^2} \gamma_{n+1} \left[3p_1 + \frac{1}{2\gamma_1^2[1 - 4\epsilon(1 + \tilde{\rho}_1)\gamma_1^2]} \right] + \frac{6(1+\tilde{\rho}_1)^2}{\pi} + 1 \\ & \leq \frac{(1+\tilde{\rho}_1)^2}{\pi} \left[24\epsilon\gamma_1^2 p_1 + \frac{4\epsilon}{1 - 4\epsilon(1 + \tilde{\rho}_1)\gamma_1^2} + 7 \right] \end{aligned}$$

and

$$B_n \left(\frac{2(1+\tilde{\rho}_1)}{\pi} \right) \geq \frac{12(1+\tilde{\rho}_1)^2}{\pi}.$$

As a consequence, $P_n(1 - \xi_n)$ is negative if we have

$$5 \geq 24\epsilon\gamma_1^2 + \frac{4\epsilon}{1 - 4\epsilon(1 + \tilde{\rho}_1)\gamma_1^2}$$

From the constraint on γ_1 , another computation shows that the above condition is fulfilled when $\epsilon^2 \frac{128(1+\tilde{\rho}_1)}{3} - \epsilon[84 + 40(1 + \tilde{\rho}_1)] + 45 \geq 0$. We then observe that all values of ϵ in $(0, \frac{1}{3}]$ can be conveniently used when $\tilde{\rho}_1 \leq \frac{227}{232}$.

To obtain *ii*), the main idea is to use the sharp estimation of the sign of P_n on $[0, 1]$ and to obtain an upper bound of $\mathbb{E}Z_n^{(3)}$. Note that:

$$\begin{aligned}
& \sup_{0 \leq t \leq 1} \gamma_{n+1}(1-t)P_n(t) \\
&= \gamma_{n+1} \sup_{1-\xi_n \leq t \leq 1} (1-t)P_n(t) \\
&= \gamma_{n+1} \sup_{1-\xi_n \leq t \leq 1} \left\{ (1-t)^3 [\epsilon_n - 3\pi t] - 3\tilde{\rho}_1(1-t)^2 \kappa_\sigma(t) \right. \\
&\quad \left. + 3[t(1-t)^3 p_1 + t^2(1-t)^2 p_2] + \gamma_{n+1} [-t(1-t)^3 p_1 + t^3(1-t)p_2] \right\}
\end{aligned}$$

We have seen in the proof of *i*) that $t \in [1-\xi_n, 1] \implies \epsilon_n \leq 3\pi t$. Hence, using $\kappa_\sigma(t) \geq -(1-\sigma p_1)t^2$, we have

$$\begin{aligned}
& \sup_{0 \leq t \leq 1} \gamma_{n+1}(1-t)P_n(t) \\
&\leq \gamma_{n+1} [3\tilde{\rho}_1(1-\sigma p_1)\xi_n^2 + 3p_1\xi_n^3 + p_2\xi_n^2 + \gamma_{n+1}\xi_n] \\
&\leq \frac{C_1(\tilde{\rho}_1, p_1, p_2, \sigma)}{\pi^2} \gamma_{n+1}^3 + \frac{C_2(\tilde{\rho}_1, p_1)}{\pi^3} \gamma_{n+1}^4
\end{aligned}$$

with $C_1(\tilde{\rho}_1, p_1, p_2, \sigma) = (1+\tilde{\rho}_1)(12\tilde{\rho}_1(1+\tilde{\rho}_1)(1-\sigma p_1) + 4p_2(1+\tilde{\rho}_1) + 2\pi)$ and $C_2(\tilde{\rho}_1, p_1) = 24p_1(1+\tilde{\rho}_1)^3$ shortened in C_1 and C_2 below. We apply Lemma 4.1 to upper bound $\sup_{n \geq n_0} \mathbb{E}Z_n^{(3)}$:

$$\begin{aligned}
& \sup_{n \geq n_0} \mathbb{E}Z_n^{(3)} \\
&\leq \mathbb{E}Z_{n_0}^{(3)} + \sup_{n \geq n_0} \mathbb{E} \sum_{k=n_0}^n \Delta Z_{n+1}^{(3)} \\
&\leq \mathbb{E}Z_{n_0}^{(3)} + \sup_{n \geq n_0} \mathbb{E} \left[\sum_{k=n_0}^n \gamma_{k+1}(1-X_k)P_k(X_k) + \Delta R_k \right] \\
&\leq \mathbb{E}Z_{n_0}^{(3)} + \frac{C_1}{\pi^2} \sum_{k=n_0}^\infty \gamma_{k+1}^3 + \frac{C_2}{\pi^3} \sum_{k=n_0}^\infty \gamma_{k+1}^4 + \sum_{k=n_0}^\infty \mathbb{E}\Delta R_k
\end{aligned}$$

Using a simple comparison argument with the integrals $\int_{n_0}^\infty t^{-\alpha} dt$, we obtain:

$$\sum_{k=n_0}^\infty \gamma_{k+1}^3 \leq 2\gamma_1^3 n_0^{-1/2} \leq 4\gamma_1^4 \epsilon \pi \quad \text{and} \quad \sum_{k=n_0}^\infty \gamma_{k+1}^4 \leq \gamma_1^4 n_0^{-1} \leq 4\gamma_1^6 \epsilon^2 \pi^2.$$

We then deduce that:

$$\sup_{n \geq n_0} \mathbb{E}Z_n^{(3)} \leq \mathbb{E}Z_{n_0}^{(3)} + \frac{4\gamma_1^4 \epsilon C_1}{\pi} + \frac{4\gamma_1^6 \epsilon^2 C_2}{\pi} + \sum_{k=n_0}^\infty \Delta R_k.$$

The result now follows using (32). \square

Explicit bound. We can now conclude the proof of Theorem 3.2.

Proof of Theorem 3.2 (b). We consider the extreme over-penalized case obtained with $\sigma = 0$. and use a power increment until $r = 3$. Recall that $n_0 := n_0(\epsilon, \pi, \gamma_1)$ is defined by (27). In particular, $\sqrt{n_0 - 1} \leq (2\epsilon\gamma_1\pi)^{-1}$ and for $i = 1, 2, 3$, $\pi\mathbb{E}[Z_{n_0}^{(i)}] \leq (2\epsilon\gamma_1^2)^{-1} + (\gamma_1)^{-1}$. Taking the results of Proposition 4.2 *ii*) and Corollary 4.1 and plugging them into (19), a series of computations yields:

$$\frac{\sup_{p_1 \geq p_2} \bar{R}_n}{\sqrt{n}} \leq c(\gamma_1, \tilde{\rho}_1, \epsilon) := T_1(\gamma_1, \tilde{\rho}_1, \epsilon) + \frac{2\gamma_1}{(1-\epsilon)(1-\epsilon/2)} T_2(\gamma_1, \tilde{\rho}_1, \epsilon), \quad (38)$$

where

$$T_1(\gamma_1, \tilde{\rho}_1, \epsilon) = \frac{1}{2\epsilon\gamma_1} + \left(\frac{1}{\epsilon\gamma_1} + 2 \right) \left(1 + \frac{1}{1-\epsilon} + \frac{1}{(1-\epsilon)(1-\epsilon/2)} \right) \\ + 2\rho_1 \left(\frac{1}{1-\epsilon} + \frac{1}{(1-\epsilon)(1-\epsilon/2)} \right) + \frac{\gamma_1}{(1-\epsilon)(1-\epsilon/2)},$$

and

$$T_2(\gamma_1, \tilde{\rho}_1, \epsilon) \\ = \gamma_1^4 \left[8\epsilon(1 + \tilde{\rho}_1) (1 + (1 + \tilde{\rho}_1)(2 + 6\tilde{\rho}_1 + 12\gamma_1^2\epsilon)) + 12\tilde{\rho}_1^2\epsilon + \frac{16}{3}\gamma_1^4\tilde{\rho}_1^3\epsilon^3 \right].$$

Theorem 3.2(b) follows by minimizing $(\gamma_1, \tilde{\rho}_1, \epsilon) \mapsto c(\gamma_1, \tilde{\rho}_1, \epsilon)$ under the constraints:

$$\epsilon \leq 1/3, \quad \frac{1}{3\sqrt{2}\tilde{\rho}_1} \leq \gamma_1^2 \leq \frac{3}{2(1 + \tilde{\rho}_1)}, \quad \tilde{\rho}_1 \leq 227/232.$$

The “best” upper bound was obtained by setting $\gamma_1 = 0.89, \tilde{\rho}_1 = 0.38, \epsilon = 1/9$, leading to the regret upper bound

$$\bar{R}_n \leq 44\sqrt{n}.$$

□

4.5 Proof of Theorems 3.1 and 3.2 (a)

We prove these results together. We thus consider $\gamma_1 \in (0, 1)$, $\rho_1 \in (0, 1)$ and $\sigma \in [0, 1]$. A variant of Proposition 4.1 concerning the increase of exponent is still valid. First, it can be observed that if we set $\varepsilon_r = r - 1/2$ (so that $\alpha_r = \pi/2$), then Lemma A.1 can be applied with $\tilde{n} \geq (\frac{\pi}{2}\gamma_1)^{-2}$. Thus, we set $n_0(\lambda) := \lfloor \frac{\lambda^2}{\pi^2} \rfloor + 1$ with $\lambda \geq 2\gamma_1^{-1}$. After a simple adaptation of the proof of Proposition 4.1, it can be deduced that for every $r \geq 1$,

$$\sup_{n \geq n_0(\lambda)} \mathbb{E}Z_n^{(r)} \leq \mathbb{E}Z_{n_0}^{(r)} + \frac{2r}{\pi} \left[\tilde{\rho}_1 + h_r(\gamma_{n_0(\lambda)}) + \pi \sup_{n \geq n_0(\lambda)} Z_n^{(r+1)} \right].$$

By an iteration, it follows by using the fact that $\pi \mathbb{E}[Z_{n_0(\lambda)}^{(i)}] \leq \pi \gamma_{n_0(\lambda)}^{-1} \leq \gamma_1^{-1}(\lambda + 1)$ that for every $r \geq 1$, some constants $C_r^1(\lambda)$ and $C_r^2(\lambda)$ exist (depending only on σ, γ_1 and ρ_1) such that,

$$\sup_{n \geq n_0(\lambda)} \pi \mathbb{E}[Y_n] \leq C_r^1(\lambda) + C_r^2(\lambda) \pi \sup_{n \geq n_0(\lambda)} \mathbb{E}Z_n^{(r+1)}. \quad (39)$$

It remains to upper bound $\sup_{n \geq n_0(\lambda)} \mathbb{E}Z_n^{(r)}$ for r large enough. Once again, a simple adaptation of the proof of Lemma 4.1 for $r \geq 3$ yields:

$$\mathbb{E}[\Delta Z_{n+1}^{(r)} | \mathcal{F}_n] = \gamma_{n+1}(1 - X_n)^{r-1} P_n^{(r)}(X_n) + \Delta R_n^{(r)}.$$

with

$$P_n^{(r)}(x) \\ = \frac{(1-x)^2}{\gamma_{n+1}} (\varepsilon_n - r\pi x) - r\tilde{\rho}_1(1-x)\kappa_\sigma(x) + \binom{r}{r-2} (x(1-x)^2 p_1 + x^2(1-x)p_2) \\ + \gamma_{n+1} \binom{r}{r-3} (-x(1-x)^2 p_1 + x^3 p_2) \quad (40)$$

and $\Delta R_n^{(r)} \leq C_r \gamma_{n+1}^3$ (where C_r does not depend on π). We want to prove that $P_n^{(r)}$ is negative on $[0, 1 - \xi_n]$ with $\xi_n = \xi \gamma_{n+1} \in (0, 1)$ where ξ is a constant to be calibrated. We follow the lines

of the proof of Proposition 4.2, but we can use some rougher arguments since we are not looking for explicit constants. First, $P_n^{(r)}(0) = \frac{\varepsilon_n}{\gamma_{n+1}} - r\tilde{\rho}_1\kappa_\sigma(0)$, so that:

$$P_n^{(r)}(0) < 0 \iff \gamma_1\rho_1 > \frac{1}{r\sqrt{2}(1-\sigma p_2)}.$$

On the one hand, for every $\sigma < 1$, it is possible to find an r sufficiently large for which this condition holds. On the other hand, when $\sigma = 1$ (case of Theorem 3.1), we then need to assume that a $\delta > 0$ exists such that $p_2 < 1 - \delta$ (in this case, the condition is satisfied if $r > (\gamma_1\rho_1\sqrt{2}\delta)^{-1}$). For such an r , it can be observed that the leading coefficient $a_n^{(r)}$ (related to x^3) is:

$$a_n^{(r)} = \left(-\frac{r}{\gamma_{n+1}} + \binom{r}{r-2} + r\sigma\tilde{\rho}_1 - \gamma_{n+1} \right) \pi.$$

It can therefore be deduced that $a_n^{(r)}$ is negative for every $n \geq n_1^\sigma$ where:

$$n_1^\sigma := \left\lceil \gamma_1^2 \left(\frac{r-1}{2} + \sigma\tilde{\rho}_1 \right)^2 \right\rceil.$$

Assume that $\lambda \geq \sqrt{n_1^\sigma}$ in order to obtain $n_0(\lambda) \geq n_1^\sigma$. Since $P_n^{(r)}(1) = \gamma_{n+1}\binom{r}{r-3}p_2 > 0$ and $\deg(P_n^{(r)}) = 3$, it follows that $P_n^{(r)}$ has exactly one root in $(0, 1)$ for every $n \geq n_0$ and that $P_n^{(r)}$ is negative on $[0, 1 - \xi_n]$ as soon as $P_n^{(r)}(1 - \xi_n) < 0$. Let n be such that $\xi_{\gamma_{n+1}} \leq 1/2$. Then, some rough estimations yield that $P_n^{(r)}(1 - \xi_n)$ is negative if

$$\frac{r\pi}{2}\xi^2 - c_r\xi - 1 > 0,$$

where c_r is a constant that does not depend on π . We then check that another constant η_r exists such that the previous property is fulfilled if $\xi \geq \eta_r/\pi$. Then, $P_n^{(r)}(1 - \frac{\eta_r}{\pi}\gamma_{n+1}) < 0$ is negative as soon as $\xi_{\gamma_{n+1}} < 1/2$. This is true for every $n \geq n_0(\lambda)$ as soon as $\lambda \geq 2\gamma_1\eta_r$. We can conclude from what precedes that an $r \geq 3$ and $\lambda > 0$ exist such that for every $n \geq n_0(\lambda)$, for every $(p_1, p_2) \in [0, 1]^2$, such that $p_1 > p_2$ (resp. $p_1 > p_2$ and $p_2 < 1 - \delta$) if $\sigma < 1$ (resp. if $\sigma = 1$)

$$\mathbb{E}[\Delta Z_{n+1}^{(r)}] \leq \gamma_{n+1} \sup_{t \in [1 - \frac{\eta_r}{\pi}\gamma_{n+1}, 1]} (1-t)P_n^{(r)}(t) + C_r\gamma_{n+1}^3.$$

Using $\gamma_{n+1} \leq \pi/\lambda$ if $n \geq n_0(\lambda)$, a constant C_λ exists such that on

$$\forall t \in [1 - \frac{\eta_r}{\pi}\gamma_{n+1}, 1] \quad P_n^{(r)}(t) \leq C_\lambda\gamma_{n+1}/\pi.$$

Under the previous conditions, we deduce

$$\sup_{\pi} \left(\pi \sup_{n \geq n_0(\lambda)} \mathbb{E}Z_n^{(r)} \right) \leq \sup_{\pi} \left(\pi \sum_{n \geq n_0} C\gamma_{n+1}^3 (\pi^{-2} + \pi^{-1}) \right) < +\infty.$$

The result follows by plugging this inequality into (39).

5 Almost sure and weak limit of the over-penalized bandit

We provide here the proofs of Propositions 3.1 and 3.2. For the sake of simplicity, we restrict our study to $\sigma = 1$ (always over-penalization of the bandit), and the argument can be adapted for any values of $\sigma \in (0, 1]$.

5.1 A.s. convergence of the multi-armed bandit (Proposition 3.1)

Recall first that $X_n = (X_{n,1}, \dots, X_{n,d})$, the multi-armed penalized bandit (13) makes it possible to define for $i \in \{2, \dots, d\}$,

$$X_{n+1,i} = X_{n,i} + \gamma_{n+1} h_i(X_n) + \gamma_{n+1} \rho_{n+1} \kappa_i(X_n) + \gamma_{n+1} \Delta M_{n+1,i},$$

where the main part of the drift h_i is defined as

$$h_i(x_1, \dots, x_d) = (1 - x_i) x_i p_i - x_i \sum_{j \neq i} x_j p_j,$$

and the penalty drift is

$$\kappa_i(x_1, \dots, x_d) = -x_i^2(1 - p_i) + \frac{1}{d-1} \sum_{j \neq i} x_j^2(1 - p_j).$$

Hence, the martingale increment is simply obtained as

$$\begin{aligned} \Delta M_{n+1,i} &= ((1 - X_{n,i}) 1_{V_{n+1,i}, A_{n+1,i}} - X_{n,i} \sum_{j \neq i} 1_{V_{n+1,j}, A_{n+1,j}} - h_i(X_n)) \\ &- \rho_{n+1} (X_{n,i} 1_{V_{n+1,i}, A_{n+1,i}^c} - \frac{1}{d-1} \sum_{j \neq i} X_{n,j} 1_{V_{n+1,j}, A_{n+1,j}^c} + \kappa_i(X_n)) \end{aligned}$$

Proof of Proposition 3.1. We start by (i) and identify the stationary points of the ODE method. The ODE $\dot{x} = h(x)$ possesses a finite number of equilibria that can be easily identified. We begin by solving the equation $h_1(x) = 0$. Since

$$h_1(x) = x_1 \sum_{i=2}^d x_i (p_1 - p_i) \geq 0,$$

we either have $x_1 = 1$ and $x_2 = \dots = x_d = 0$ or $x_1 = 0$.

Then, the equation $h_2(x) = 0$ with $x_1 = 0$ may be reduced to

$$x_2 \sum_{i=3}^d x_i (p_2 - p_i) \geq 0.$$

The same argument leads to $x_2 = 1$ or $x_2 = 0$ and a straightforward recursion shows that the equilibria of the ODE are $(\delta^i)_{1 \leq i \leq d}$, with $(\delta^i)_{1 \leq i \leq d}$ defined as

$$\delta_i^i = 1 \quad \text{and} \quad \delta_j^i = 0 \quad \forall j \neq i.$$

Let us emphasize that to discriminate among these equilibria, it is not possible to use the second derivative criterion that relies on $\left(\frac{\partial h_i}{\partial x_j}\right)_{i,j}$ to establish their stability. Instead, it is possible to check that δ^1 fulfills the Lyapunov certificate with the function $V(x) = (x_2^2 + \dots + x_d^2)$. If we denote $h = (h_1, \dots, h_d)$, we then have:

$$\langle \nabla V(x), h(x) \rangle = \sum_{j=2}^d x_j^2 \sum_{k \neq j} x_k (p_j - p_k).$$

Considering x in a closed neighborhood of δ^1 defined as $x_j \leq \epsilon/d$, $\forall j \geq 2$ (implying that $x_1 > 1 - \epsilon$), we see that:

$$\begin{aligned} \langle \nabla V(x), h(x) \rangle &= x_1 \sum_{j=2}^d x_j^2 (p_j - p_1) + \sum_{k=2}^d x_k^2 \sum_{j \neq k, j \geq 2} x_j (p_k - p_j) \\ &\leq -(1 - \epsilon)(p_1 - p_2) \sum_{j=2}^d x_j^2 + \epsilon \sum_{j=2}^d x_j^2, \end{aligned}$$

and the term above is negative as soon as ϵ is chosen such that:

$$\epsilon \leq \frac{1}{p_1 - p_2 + 1}.$$

In contrast, the other equilibria $(\delta^j), j \neq 1$ are unstable: this can be easily deduced from the instability of the two-armed bandit by testing the first arm vs. the arm j .

Since the martingale increment $\Delta M_{n+1,i}$ is uniformly bounded, we can apply the Kushner-Clark theorem (see [14]) and can conclude that $(X_{n,i})_{n \geq 0}$ either converges to 1 or 0 a.s. As a consequence, it is also true that $(X_n)_{n \geq 0}$ converges a.s. We now make this limit explicit and show that $(X_n)_{n \geq 0}$ converges toward $(1, \dots, 0)$ a.s. We start by noticing that $h_1(x) = x_1 \sum_{j \geq 2} x_j (p - 1 - p_j) \geq 0$, which implies that:

$$X_{n,1} \geq X_{0,1} + \sum_{j=1}^{n-1} \gamma_j \rho_j \kappa_{j-1,1}(X_{j-1}) + \sum_{j=1}^{n-1} \gamma_j \Delta M_j. \quad (41)$$

The martingale increment ΔM_j is bounded and a large enough C exists such that $\Delta M_j \leq \sqrt{C}$. This implies that:

$$\left\| \sum_{j=1}^{n-1} \gamma_j \Delta M_j \right\|_{L^2}^2 \leq C \sum_{j=1}^{n-1} \gamma_j^2 \leq C \sup_{j \in \mathbb{N}} \left(\frac{\gamma_j}{\rho_j} \right) \sum_{j=1}^{n-1} \gamma_j \rho_j.$$

Since $\sum \rho_j \gamma_j = +\infty$, we can deduce that

$$\lim_{n \rightarrow +\infty} \frac{\mathbb{E} \left[\sum_{j=1}^n \gamma_j \Delta M_j \right]^2}{\sum_{j=1}^{n-1} \gamma_j \rho_j} = 0 \quad \text{so that} \quad \limsup_{n \rightarrow \infty} \frac{\sum_{j=1}^{n-1} \gamma_j \Delta M_j}{\sum_{j=1}^{n-1} \gamma_j \rho_j} \geq 0.$$

We now consider an event $\omega \in \{X_{\infty,1} = 0\}$. We have:

$$\lim_{n \rightarrow +\infty} \kappa_1(X_n(\omega)) = \frac{1}{d-1} \sum_{k \geq 2} (1 - p_k) X_{\infty,k}(\omega)^2,$$

and according to the Toeplitz Lemma we deduce that

$$\lim_{n \rightarrow \infty} \frac{\sum_{j=1}^{n-1} \gamma_j \rho_j \kappa_1(X_{j-1}(\omega))}{\sum_{j=1}^{n-1} \gamma_j \rho_j} = \frac{1}{d-1} \sum_{k \geq 2} (1 - p_k) X_{\infty,k}(\omega)^2 > 0.$$

Putting together this last remark with Equation (41) leads to the conclusion

$$\limsup_{n \rightarrow \infty} \frac{X_{n,1}(\omega)}{\sum_{j=1}^{n-1} \gamma_j \rho_j} > 0.$$

We obtain a contradiction with the boundedness of $(X_n)_{n \geq 1}$ and conclude that $\mathbb{P}(X_{\infty,1} = 0) = 0$. For (ii), we refer to [16] since the arguments here are similar. \square

5.2 Weak convergence of the normalized bandit (Proposition 3.2)

The proof of the weak convergence follows the lines of [16]. The idea is to prove the tightness of the pseudo-trajectories associated to the normalized sequence and to then show that any weak limit of this sequence is a solution of the martingale problem $(\mathcal{L}, \mathcal{C}_K^1(\mathbb{R}_+, (\mathbb{R}_+)^{d-1})$ where \mathcal{L} is the infinitesimal generator defined in Proposition 3.2. Then, proving that uniqueness holds for the solutions of the martingale problem and for the invariant distribution, the convergence follows. Here, we choose to only detail the key step of the characterization of the limit. The rest of the proof can be obtained by a simple generalization of that of [16].

Proposition 5.1. *Let f be a continuously differentiable function with compact support in \mathbb{R}_+^{d-1} . We have*

$$\mathbb{E}(f(Y_{n+1,2}, \dots, Y_{n+1,d}) - f(Y_{n,2}, \dots, Y_{n,d}) | \mathcal{F}_n) = \gamma_{n+1} \mathcal{L}_d f(Y_{n,2}, \dots, Y_{n,d}) + o_P(1),$$

where \mathcal{L}_d is the PDMP generator defined in (15) and $\mathcal{F}_n = \sigma(Y_k, k \leq n)$.

Proof. Since the proof does not depend on σ , we assume that $\sigma = 1$ for the sake of clarity. We first give an alternative expression for the variables $Y_{n,i}$ for $i \geq 2$.

$$Y_{n+1,i} = Y_{n,i} + \gamma_{n+1} \left(\frac{1-p_1}{d-1} - (p_1 - p_i)Y_{n,i} \right) + \gamma_{n+1} C_{n,i} - g \Delta M_{n+1,i},$$

where $C_{n,i} = (\kappa_i(X_n) - \frac{1-p_1}{d-1}) + Y_{n,i}(p_1 - p_i + (\epsilon_n + \frac{\rho_n}{\rho_{n+1}}(p_i - \sum_{j \neq i} X_{n,j} p_j))) = o_P(1)$ since $(\epsilon_n)_{n \geq 0}$ converges 0 and $(X_{n,i})_{n \geq 0}$ converges to 0 in probability for $i \geq 2$. We rewrite this as follows

$$Y_{n+1,i} = Y_{n,i} + \gamma_{n+1} \left(\frac{1-p_1}{d-1} - (p_1 - p_i)Y_{n,i} + C_{n,i} \right) + G_{n,i} + g \Delta \tilde{M}_{n+1,i},$$

where $G_{n,i} = g(1 - X_{n,i})(1_{V_{n+1,i}, A_{n+1,i}} - X_{n,i} p_i)$ and $\Delta \tilde{M}_{n+1,i} = \Delta M_{n+1,i} - G_{n,i}$.

We consider a function $f \in \mathcal{C}^1(\mathbb{R}_+^{d-1})$ with a compact support.

$$f(Y_{n+1}) - f(Y_n) = \sum_{i=2}^d f(Y_{n+1,2}, \dots, Y_{n+1,i}, \dots, Y_{n+1,d}) - f(Y_{n,2}, \dots, Y_{n,i}, \dots, Y_{n,d}).$$

We will use the following notation $F_i(Y_k) = f(Y_{k,2}, \dots, Y_{k,i}, \dots, Y_{k,d})$. This means that the first $i-1$ variables are $(Y_{n,2}, Y_{n,3}, \dots)$ and the last $d-i$ ones are: $(Y_{n+1,i+1}, Y_{n+1,i+2}, \dots, Y_{n+1,d})$. We have:

$$F_i(Y_{n+1,i}) - F_i(Y_{n,i}) = F_i(Y_{n+1,i}) - F_i(\bar{Y}_{n,i}) + F_i(\bar{Y}_{n,i}) - F_i(Y_{n,i}),$$

where

$$\tilde{Y}_{n,i} = Y_{n,i} + \gamma_{n+1} \left(\frac{1-p_1}{d-1} - (p_1 - p_i)Y_{n,i} + C_{n,i} \right),$$

and

$$\bar{Y}_{n,i} = \tilde{Y}_{n,i} + G_{n,i}.$$

We begin by writing:

$$F_i(Y_{n+1,i}) - F_i(\bar{Y}_{n,i}) = \partial_i F_i(\tilde{Y}_{n,i}) \Delta \tilde{M}_{n+1,i} + \gamma_{n+1} V_{n+1,i},$$

where the first order Taylor approximation formula yields:

$$\exists \theta \in [0, 1] : \quad V_{n+1,i} = \left[F_i(\tilde{Y}_{n,i} + \theta \Delta \tilde{M}_{n+1,i}) - F_i(\tilde{Y}_{n,i}) \right] \Delta \tilde{M}_{n+1,i}.$$

As a consequence, $V_{n+1,i} = o_P(1)$ and we are now going to prove that:

$$\mathbb{P} - \lim_{n \rightarrow \infty} \mathbb{E} \left(\frac{F_i(Y_{n+1}) - F_i(Y_n) - \gamma_{n+1} \mathcal{L}_d F_i(Y_n)}{\gamma_{n+1}} | \mathcal{F}_n \right) = 0,$$

where:

$$\begin{aligned}\mathcal{A}_i f(Y_2, \dots, Y_d) &= \frac{p_i Y_i}{g} (f(Y_2, \dots, Y_i + g, \dots, Y_d) - f(Y_2, \dots, Y_i, \dots, Y_d)) \\ &+ \left(\frac{1-p_1}{d-1} - p_1 Y_i \right) \partial_i f(Y_2, \dots, Y_d).\end{aligned}$$

We compute:

$$\begin{aligned}\mathbb{E}(F_i(\bar{Y}_{n,i}) | \mathcal{F}_{n,i}) &= p_i X_{n,i} F_i(\tilde{Y}_{n,i} + g(1 - X_{n,i})(1 - p_i X_{n,i})) \\ &+ (1 - gp_i X_{n,i}) F_i(\tilde{Y}_{n,i} - gp_i X_{n,i}(1 - X_{n,i})).\end{aligned}$$

Let us decompose the r.h.s. of the above equation into two parts, denoted by:

$$F_{n,i} = p_i X_{n,i} (F_i(\tilde{Y}_{n,i} + g(1 - X_{n,i})(1 - p_i X_{n,i})) - F_i(Y_{n,i})), \quad (42)$$

and

$$G_{n,i} = (1 - gp_i X_{n,i}) (F_i(\tilde{Y}_{n,i} - gp_i X_{n,i}(1 - X_{n,i})) - F_i(Y_{n,i})). \quad (43)$$

Note that (42) is the jump part of the PDMP and (43) is the deterministic one. If $i \geq 2$, $(X_{n,i})_{n \geq 1}$ converges to 0 in probability and $\rho_n \gamma_{n+1}^{-1} = g + o(\rho_n)$. Thus:

$$\begin{aligned}\gamma_{n+1}^{-1} F_{n,i} &= \gamma_{n+1}^{-1} \rho_n p_i Y_{n,i} (F_i(Y_{n,i} + g + o_P(1)) - F_i(Y_{n,i})) \\ &= \frac{p_i Y_{n,i}}{g} (1 + o(\rho_n)) [F_i(Y_{n,i} + g + o_P(1)) - F_i(Y_{n,i})].\end{aligned}$$

As a consequence, the asymptotic behavior of (42) is given by

$$\mathbb{P} - \lim_{n \rightarrow \infty} \left(\frac{F_{n,i}}{\gamma_{n+1}} - p_i Y_{n,i} \frac{F_i(Y_{n,i} + g) - F_i(Y_{n,i})}{g} \right) = 0.$$

We now study (43) and compute:

$$\begin{aligned}\tilde{Y}_{n,i} - g X_{n,i} (1 - p_i X_{n,i}) &= Y_{n,i} + \gamma_{n+1} \left(\frac{1-p_1}{d-1} - p_1 Y_{n,i} \right) \\ &+ \gamma_{n+1} p_i Y_{n,i} - gp_i X_{n,i} (1 - X_{n,i}) + \gamma_{n+1} C_{n,i} \\ &= Y_{n,i} + \gamma_{n+1} \left(\frac{1-p_1}{d-1} - p_1 Y_{n,i} \right) \\ &+ \underbrace{\gamma_{n+1} p_i Y_{n,i} - gp_i X_{n,i} (1 - X_{n,i}) + \gamma_{n+1} C_{n,i}}_{:= \gamma_{n+1} \tilde{C}_{n,i}},\end{aligned}$$

where $g\rho_n = \gamma_n$. Since $\tilde{C}_{n,i}$ converges to 0 in probability, we obtain:

$$\begin{aligned}\gamma_{n+1}^{-1} G_{n,i} &= \gamma_{n+1}^{-1} (1 + o(\rho_n)) \left(F_i(Y_{n,i} + \gamma_{n+1} \left[\frac{1-p_1}{d-1} - p_1 Y_{n,i} \right] + \gamma_{n+1} \tilde{C}_{n,i}) - F_i(Y_{n,i}) \right) \\ &= \left[\frac{1-p_1}{d-1} - p_1 Y_{n,i} \right] \frac{\left(F_i(Y_{n,i} + \gamma_{n+1} \left[\frac{1-p_1}{d-1} - p_1 Y_{n,i} \right] + \gamma_{n+1} \tilde{C}_{n,i}) - F_i(Y_{n,i}) \right)}{\gamma_{n+1} \left[\frac{1-p_1}{d-1} - p_1 Y_{n,i} \right]} \\ &+ o_P(1).\end{aligned}$$

We finally obtain the limiting behavior of (43):

$$\mathbb{P} - \lim_{n \rightarrow \infty} \left(\frac{G_{n,i}}{\gamma_{n+1}} - \left(\frac{1-p_1}{d-1} - p_1 Y_i \right) \partial_i F_i(Y_{n,i}) \right) = 0.$$

This ends the proof of the proposition. \square

6 Ergodicity of the PDMP

From now on, the variable $(X_t)_{t \geq 0}$ will refer to a trajectory of the PDMP associated with the normalized (over)-penalized bandit and bearing no relation to the multi-armed bandit sequence $(X_n)_{n \geq 1}$.

6.1 Wasserstein results

We begin the study of the ergodicity of the PDMP whose infinitesimal generator is (16) with some computations of the moments of the process.

Lemma 6.1. *Let $(X_t)_{t \geq 0}$ be a Markov process, whose generator \mathcal{L} is defined by (16). If $\pi := b - cg > 0$, then $\sup \mathbb{E}[(X_t^x)^p] \leq C(1 + |x|^p)$. In particular, the invariant distribution π has moments of any order and*

$$\forall t \geq 0 \quad \mathbb{E}(X_t) = \frac{a}{\pi} + \left(\mathbb{E}(X_0) - \frac{a}{\pi} \right) e^{-t\pi}$$

Proof. Let us define $f_p(x) = x^p$. We have:

$$\begin{aligned} \mathcal{L}f_p(x) &= p(a - bx)x^{p-1} + cx((x+g)^p - x^p) \\ &= -p\pi f_p(x) + pa f_{p-1}(x) + c \sum_{k=0}^{p-2} C_p^k g^{p-k} f_{k+1}(x), \end{aligned} \quad (44)$$

where we adopt the convention $\Sigma_\emptyset = 0$. If we now define $\alpha_p(t) = \mathbb{E}(X_t^p)$, the previous relation shows that α_p satisfies the ODE for any integer $p \geq 1$ defined by

$$\alpha_p(t)' + p\pi\alpha_p(t) = pa\alpha_{p-1}(t) + c \sum_{k=0}^{p-2} C_p^k g^{p-k} \alpha_{k+1}(t).$$

For example, with $p = 1$ we have $\alpha_1'(t) = -\pi\alpha_1(t) + a$, which implies that

$$\alpha_1(t) = \frac{a}{\pi} + \left(\mathbb{E}(X_0) - \frac{a}{\pi} \right) e^{-t\pi}.$$

The control of the moments of order $p > 1$ then follows from a recursion. \square

6.1.1 Rescaled two-armed bandit & Theorem 3.3

In the following, we will exploit Equation (44) to obtain a suitable upper bound of the Wasserstein distance \mathcal{W}_p between the law of X_t and the invariant measure μ_∞ of the PDMP. For this purpose, we note that the generator (16) possesses the stochastic monotonicity property, *i.e.*, a coupling (X, Y) exists starting from (x, y) (with $x > y$) such that $X_t \geq Y_t$ for any $t \geq 0$. The increase of the jump rate (with respect to the position) and the positivity of the jumps are of prime importance for this property. Such a coupling could be built as follows: we only allow simultaneous jumps of both components or a single jump of the highest one (see ([7]) for a similar procedure). The generator of this coupling (X, Y) starting from (x, y) with $x > y$ is given by:

$$\begin{aligned} \mathcal{L}_{\mathcal{W}}f(x, y) &= (a - bx)\partial_x f(x, y) + (a - by)\partial_y f(x, y) \\ &\quad + cy(f(x + g, y + g) - f(x, y)) + c(x - y)(f(x + g, y) - f(x, y)) \end{aligned} \quad (45)$$

with a symmetric expression when $y > x$. We now prove the main result.

Proof of theorem 3.3. Let μ_0 be a probability on \mathbb{R}_+^* and designate μ_∞ as the invariant distribution of the PDMP. Set

$$\mathcal{C}_t = \{\nu \in \mathcal{P}(\mathbb{R}^2), \nu(dx \times \mathbb{R}_+) = \mu_t(dx), \nu(\mathbb{R}_+ \times dy) = \mu_\infty(dy)\}.$$

For any $\nu \in \mathcal{C}$, let $(X_t, Y_t)_{t \geq 0}$ denote the Markov process driven by (45) starting from ν . From the definition of \mathcal{W}_p and the stationary of (Y_t) , we have for any t :

$$\mathcal{W}_p(\mu_t, \mu_\infty) \leq \inf\{\nu \in \mathcal{C}_0, \left(\int_{\mathbb{R}_+^2} \mathbb{E}[|X_t^x - Y_t^y|^p] \nu(dx, dy) \right)^{\frac{1}{p}} \}.$$

At the price of a potential exchange of the coordinates, we can now work with some deterministic starting points x and y such that $x > y > 0$. Owing to the monotonicity of $\mathcal{L}_\mathcal{W}$, we thus have for any $p \geq 1$

$$\mathbb{E}(|X_t^x - Y_t^y|)^p = \mathbb{E}(X_t^x - Y_t^y)^p.$$

Assume now that $p \in \mathbb{N}^*$, we observe that $\mathcal{L}_\mathcal{W}$ acts on $(x, y) \mapsto (x - y)^p$ as:

$$\mathcal{L}_\mathcal{W}(x - y)^p = -p\pi(x - y)^p + pa(x - y)^{p-1} + c \sum_{k=0}^{p-2} C_p^k g^{p-k} (x - y)^{k+1}.$$

Setting $\beta_p(t) = \mathbb{E}|X_t^x - Y_t^y|^p$, we can immediately check that:

$$\dot{\beta}_p(t) + \pi p \beta_p(t) = \left(pa\beta_{p-1}(t) + c \sum_{k=0}^{p-2} C_p^k g^{p-k} \beta_{k+1}(t) \right). \quad (46)$$

When $p = 1$, (46) implies that: $\beta_1(t) = \beta_1(0)e^{-\pi t} \Rightarrow \mathbb{E}[X_t^x - Y_t^y] = (x - y)e^{-\pi t}$, so that:

$$\mathcal{W}_1(\mu_t, \mu_\infty) \leq \mathcal{W}_1(\mu_0, \mu_\infty)e^{-\pi t}.$$

For the lower-bound, we use:

$$\mathcal{W}_1(\mu_t, \mu_\infty) \geq \inf \left\{ \nu_t \in \mathcal{C}_t, \left| \int (x - y) \nu_t(dx, dy) \right| \right\} = |\mathbb{E}[X_t^{\mu_0}] - \mathbb{E}[Y_t^{\mu_\infty}]|,$$

which implies that:

$$\mathcal{W}_1(\mu_t, \mu_\infty) \geq \left| \int \mathbb{E}[X_t^x - Y_t^y] \mu_0(dx) \mu_\infty(dy) \right| = \left| \int (x - y) \mu_0(dx) \mu_\infty(dy) \right| e^{-\pi t}.$$

The lower-bound follows.

Now, let us consider the case $p > 1$ (with $p \in \mathbb{N}$). For $p = 2$, we have

$$(\beta_2(t)e^{2\pi t})' e^{-2\pi t} = (2a + cg^2)\beta_1(0)e^{-\pi t},$$

and an integration leads to $\beta_2(t)e^{2\pi t} - \beta_2(0) = \frac{2a+cg^2}{\pi}\beta_1(0)[e^{\pi t} - 1]$. As a consequence:

$$\beta_2(t) \leq e^{-2\pi t}\beta_2(0) + \frac{2a + cg^2}{\pi}\beta_1(0)e^{-\pi t}.$$

Using the inequalities $\sqrt{u+v} \leq \sqrt{u} + \sqrt{v}$ and $\beta_2 \geq \mathcal{W}_2^2$, we thus deduce that:

$$\mathcal{W}_2(\mu_t, \mu_\infty) \leq \mathcal{W}_2(\mu_0, \mu_\infty)e^{-\pi t} + \sqrt{\frac{2a + cg^2}{\pi}} \sqrt{\mathcal{W}_1(\mu_0, \mu_\infty)} e^{-\frac{\pi t}{2}}.$$

The result follows when $p = 2$ by setting:

$$\gamma_2 := \mathcal{W}_2(\mu_0, \mu_\infty) + \sqrt{\frac{2a + cg^2}{\pi}} \sqrt{\mathcal{W}_1(\mu_0, \mu_\infty)}.$$

A recursive argument based on (46) shows that a constant γ_p exists that only depends on μ_0 and μ_∞ such that:

$$\mathcal{W}_p(\mu_t, \mu_\infty) \leq \gamma_p e^{-\frac{\pi}{p}t}.$$

□

6.2 Proof of total variation results

As mentioned before, the idea is to wait until the paths get close (with a probability controlled by the Wasserstein bound) and then to try to stick them (with high probability). Since the jump size is deterministic, sticking the paths implies a non trivial coupling of the jump times which is described in the lemma below.

We begin by establishing the next useful lemma.

Lemma 6.2. *Let $\varepsilon > 0$ and $t \geq \frac{1}{b} \ln(1 + \varepsilon)$. A coupling $(X_t, Y_t)_{t \geq 0}$ of paths driven by (16) exists such that on $A_{x_0, \varepsilon}$:*

$$\mathbb{P}(X_t = Y_t, t \geq s) \geq \left(1 - \frac{c}{b} x_0 \varepsilon - e^{-\frac{c}{b} c s} - \frac{c \varepsilon}{b}\right) \max(0, 1 - \frac{c}{b} \varepsilon (x_0 + g)),$$

where $A_{x_0, \varepsilon} = \{(x, y) | \frac{a}{b} < x \leq x_0, 0 < x - y \leq \varepsilon\}$.

Proof Let $\varepsilon > 0$ and $(x, y) \in A_{x_0, \varepsilon}$ (in particular, $x > y$). Designate T_1^x and T_1^y as the first jumps of (X_t^x) and (X_t^y) , respectively, and T_2^x as the second jump of (X_t^x) . It can be noted that:

$$\mathbb{P}(X_t = Y_t, t \geq s) \geq \mathbb{P}(X_{T_1^y}^x = X_{T_1^y}^y, T_1^y \leq s).$$

We aim to build a coupling that leads to a sharp lower-bound of the r.h.s. For this purpose, note that if $T_1^x < T_1^y < T_2^x$, the triple (T_1^x, T_1^y, T_2^x) satisfies:

$$X_{T_1^y}^y = X_{T_1^y}^x \iff \frac{a}{b} + \left(y - \frac{a}{b}\right) e^{-b T_1^y} + g = \frac{a}{b} + \left(X_{T_1^x}^x - \frac{a}{b}\right) e^{-b(T_1^y - T_1^x)}.$$

Considering that $X_{T_1^x}^x = \frac{a}{b} + (x - \frac{a}{b}) e^{-b T_1^x} + g$ and defining $\psi(t) = \frac{1}{b} \ln \left(e^{bt} + \frac{x-y}{g} \right)$, we can verify that $X_{T_1^y}^y = X_{T_1^y}^x \leq s$ and $T_1^x < T_1^y < T_2^x$ as soon as

$$T_1^y = \psi(T_1^x) \leq s \quad \text{and} \quad T_2^x \geq \psi(T_1^x),$$

since $\psi(t) \geq t$. We are naturally encouraged to consider $S_1^{x,s} = \psi(T_1^x) 1_{\{\psi(T_1^x) \leq s\}}$ and it is well known that the law of (T_1^x, T_1^y) can be described through the maximal coupling:

$$T_1^y = \Theta U + (1 - \Theta) V_y, \quad \psi(T_1^x) = \Theta U + (1 - \Theta) V_x,$$

where V_x, V_y, Θ and U are independent, $U \sim \frac{\mathbb{P}_{T_1^y} \wedge \mathbb{P}_{\psi(T_1^x)}}{\|\mathbb{P}_{S_1^{x,s}} \wedge \mathbb{P}_{T_1^y}\|_{TV}}$ and $\Theta \sim \mathcal{B}(p)$ where $p = \|\mathbb{P}_{S_1^{x,s}} \wedge \mathbb{P}_{T_1^y}\|_{TV}$. With this coupling, if $q(t, z) = \mathbb{P}(T_1^z \geq \psi(t) - t)$, the Strong Markov property yields

$$\mathbb{P}(T_2^x - T_1^x | (T_1^x, T_1^y)) = \mathbb{P}(T_2^x \geq \psi(T_1^x) | T_1^x) = q(T_1^x, X_{T_1^x}^x).$$

Since $z \mapsto q(t, z)$ is increasing and $x > a/b$ (from the assumption on $A_{x_0, \varepsilon}$), we deduce that $X_{T_1^x}^x \leq x + g$ and it therefore follows that:

$$\mathbb{P}(T_2^x > T_1^y | (T_1^x, T_1^y)) \geq q(t, x + g) \geq q(0, x + g).$$

given that $t \mapsto \psi(t) - t$ is a non-decreasing function. As a consequence, we obtain that with this coupling:

$$\mathbb{P}(X_{T_1^y}^x = X_{T_1^y}^y, T_1^y \leq s) \geq q(0, x + g) \mathbb{P}(\Theta = 1) = q(0, x + g) \|\mathbb{P}_{S_1^{x,s}} \wedge \mathbb{P}_{T_1^y}\|_{TV}. \quad (47)$$

It remains to find a lower bound of the total variation distance involved in the r.h.s. of the above inequality. Recall that

$$\|\mathbb{P}_{S_1^{x,s}} \wedge \mathbb{P}_{T_1^y}\|_{TV} = \int_0^{+\infty} f_y(t) \wedge g_{x,s}(t) dt,$$

where f_y and $g_{x,s}$ denote the densities of T_1^y and $S_1^{x,s}$, respectively. We therefore have:

$$\forall t > 0, \quad f_y(t) = c\phi(y, t)e^{-\int_0^t c\phi(y, u)du} \quad \text{with} \quad \phi(y, t) = \frac{a}{b} + (y - \frac{a}{b})e^{-bt},$$

and a change of variable yields:

$$\forall t > 0, \quad g_x(t) = f_x(\psi^{-1}(t))(\psi^{-1})'(t)1_{\{\psi(0) \leq t \leq s\}}. \quad (48)$$

On the one hand, since $(x, y) \in A_{x_0, \varepsilon}$, we can check that:

$$\forall t \geq 0, \quad \phi(x, t) - \varepsilon e^{-bt} \leq \phi(y, t) \leq \phi(x, t),$$

and we can then conclude that:

$$\forall t > 0, \quad f_y(t) \geq f_x(t) - \varepsilon e^{-bt}.$$

On the other hand, note that:

$$\forall t > \psi(0), \quad \psi^{-1}(t) = \frac{1}{b} \ln \left(e^{bt} - \frac{x-y}{g} \right) \leq t \quad \text{and} \quad (\psi^{-1})'(t) = \frac{e^{bt}}{e^{bt} - \frac{x-y}{g}} \geq 1,$$

and we can deduce from (48) that $\forall t \in [\psi(0), s]$:

$$g_x(t) \geq c\phi(x, \psi^{-1}(t))e^{-\int_0^t c\phi(x, s)ds} \geq c\phi(x, t)e^{-\int_0^t c\phi(x, s)ds} = f_x(t).$$

Note that we used that $t \mapsto \phi(x, t)$ is decreasing since $x > a/b$. Thus,

$$\left(\mathbb{P}_{T_1^y} \wedge \mathbb{P}_{S_1^x} \right) (dt) \geq h(t)dt \quad \text{with} \quad h(t) = (f_x(t) - \varepsilon e^{-bt})1_{\psi(0) \leq t \leq s}dt.$$

As a consequence,

$$\|\mathbb{P}_{S_1^{x,s}} \wedge \mathbb{P}_{T_1^y}\|_{TV} \geq e^{-\int_0^{\psi(0)} c\phi(x, u)du} - e^{-\int_0^s c\phi(x, u)du} - \frac{\varepsilon}{b}.$$

Checking that $\psi(0) \leq \varepsilon/b$ and that $\forall t \geq 0, a/b \leq \phi(x, t) \leq x \leq x_0$, we deduce that

$$\|\mathbb{P}_{S_1^{x,s}} \wedge \mathbb{P}_{T_1^y}\|_{TV} \geq e^{-\frac{cx_0\varepsilon}{b}} - e^{-\frac{a}{b}cs} - \frac{\varepsilon}{b} \geq 1 - \frac{cx_0\varepsilon}{b} - e^{-\frac{a}{b}cs} - \frac{\varepsilon}{b},$$

where we used $e^{-u} \geq 1 - u$ for $u \geq 0$ in the second line. To conclude the proof, it remains to plug this inequality into (47) and to observe that:

$$q(0, x+g) \geq q(0, x_0+g) = e^{-\int_0^{\psi(0)} c\phi(x_0+g, s)ds} \geq 1 - c\psi(0)(x_0+g) \geq 1 - \frac{c}{b}\varepsilon(x_0+g).$$

□

We now provide the proof of the ergodicity w.r.t. the total variation distance.

Proof of Theorem 3.4. For any starting distribution μ_0 ,

$$\|\mu_0 P_t - \mu_\infty\|_{TV} \leq \int \|\delta_x P_t - \delta_y P_t\|_{TV} \mu_0(dy) \mu_\infty(dx). \quad (49)$$

The idea is to use the Wasserstein coupling during a time t_1 and to then try to stick the paths on the interval $[t_1, t]$ using Lemma 6.2. Consider $A_{x_0, \varepsilon}$ defined in Lemma 6.2 and the alternative set $A_{x_0, \varepsilon}^* = \{(x, y), a/b < y < x_0, 0 < y - x \leq \varepsilon\}$. Set $B_{x_0, \varepsilon} = A_{x_0, \varepsilon} \cup A_{x_0, \varepsilon}^*$, we have:

$$1 - \|\delta_x P_t - \delta_y P_t\|_{TV} \geq \mathbb{P}(X_t^x = Y_t^y | (X_{t_1}^x, Y_{t_1}^y) \in B_{x_0, \varepsilon}) \mathbb{P}((X_{t_1}^x, Y_{t_1}^y) \in B_{x_0, \varepsilon}). \quad (50)$$

Since the Wasserstein coupling preserves the order and since $x > a/b \mu_\infty(dx)$ -a.s., it can be noted that $\mu_\infty(dx)$ -a.s.,

$$(X_{t_1}^x, Y_{t_1}^y) \in B_{x_0, \varepsilon} \iff \begin{cases} X_{t_1}^x - X_{t_1}^y \leq \varepsilon \text{ and } X_{t_1}^x \leq x_0 & \text{if } x \geq y \\ X_{t_1}^y - X_{t_1}^x \leq \varepsilon \text{ and } X_{t_1}^y \leq x_0 & \text{if } x < y. \end{cases}$$

It follows that for every $p > 0$, $\mu_\infty(dx)$ almost surely:

$$\begin{aligned} \mathbb{P}((X_{t_1}^x, Y_{t_1}^y) \in B_{x_0, \varepsilon}^c) &\leq \mathbb{P}(|X_{t_1}^x - X_{t_1}^y| > \varepsilon) + \mathbb{P}(X_{t_1}^x > x_0) + \mathbb{P}(X_{t_1}^y > x_0) \\ &\leq \frac{1}{\varepsilon} \mathbb{E}[|X_{t_1}^x - X_{t_1}^y|] + \frac{1}{x_0^p} (\mathbb{E}[(X_{t_1}^x)^p] + \mathbb{E}[(X_{t_1}^y)^p]). \end{aligned}$$

On the basis of Theorem 3.3 and Lemma 6.1, a constant C_p exists such that C_p depends on p , μ_0 and μ_∞ but not on t_1 and satisfies:

$$\int \mathbb{P}((X_{t_1}^x, Y_{t_1}^y) \in B_{x_0, \varepsilon}^c) \mu_0(dy) \mu_\infty(dx) \leq \frac{\mathcal{W}_1(\mu_0, \mu_\infty)}{\varepsilon} e^{-\pi t_1} + \frac{C_p}{x_0^p}.$$

Finally, Lemma 6.2 leads to:

$$\begin{aligned} &\mathbb{P}(X_t^x = Y_t^y | (X_{t_1}^x, Y_{t_1}^y) \in B_{x_0, \varepsilon}) \\ &\geq \left(1 - \frac{c}{b} x_0 \varepsilon - e^{-\frac{a}{b} c(t-t_1)} - \frac{c\varepsilon}{b}\right) \left\{0 \vee 1 - \frac{c}{b} \varepsilon (x_0 + g)\right\} \end{aligned}$$

so that by plugging the previous inequalities into (50) and (49), it can be deduced that for every $p > 1$, a constant \tilde{C}_p exists such that for every $t \geq 0$, for every x_0 and ε such that $x_0 \varepsilon \leq b/2c$ (with $x_0 > 1$ and $\varepsilon \in (0, 1)$),

$$\|\mu_0 P_t - \mu_\infty\|_{TV} \leq \tilde{C}_p \left(x_0 \varepsilon + e^{-\frac{a}{b} c(t-t_1)} + \varepsilon + \frac{1}{\varepsilon} e^{-\pi t_1} + \frac{1}{x_0^p} \right).$$

If we try to optimize the above bound, we set $t_1 = \delta t$, $x_0 = C_1 e^{\alpha t}$, $\varepsilon = C_2 e^{-\beta t}$ with $\delta \in (0, 1)$ and $\beta > \alpha > 0$ and deduce that a constant \check{C}_p exists such that:

$$\|\mu_0 P_t - \mu_\infty\|_{TV} \leq \check{C}_p \exp \left(-t \left\{ \beta - \alpha \wedge \frac{ca}{b} (1 - \delta) \wedge \delta \pi - \beta \wedge \alpha p \right\} \right).$$

We can choose p as large as we want (μ_0 has moments of any order) and thus α arbitrarily small. The result then follows using an optimization on (β, δ) . \square

A Technical result for the pseudo-regret upper bound

Lemma A.1. *Let $\alpha > 0$, $\gamma_1 \in (0, 1)$ and $\tilde{n} \in \mathbb{N}$ such that $\alpha \gamma_{\tilde{n}} < 1$ and $\tilde{n} \geq 1/(\alpha \gamma_1)^2$. We have:*

$$\forall n \geq \tilde{n} \quad \sum_{j=\tilde{n}}^{n-1} \gamma_j \prod_{l=j}^{n-1} (1 - \alpha \gamma_l) \leq \frac{1}{\alpha}$$

Proof Let $j \geq \tilde{n}$. On the basis of the inequality $\ln(1+x) \geq x$ for $x > -1$, we have

$$\prod_{l=j}^{n-1} (1 - \alpha \gamma_l) = \exp \left(\ln \sum_{l=j}^{n-1} (1 - \alpha \gamma_l) \right) \leq \exp \left(- \sum_{l=j}^{n-1} \alpha \gamma_l \right)$$

Using that $x \mapsto 1/\sqrt{x}$ is decreasing,

$$\sum_{l=j}^{n-1} \gamma_l = \sum_{l=j}^{n-1} \frac{\gamma_1}{\sqrt{l}} \geq \gamma_1 \sum_{l=j}^{n-1} \int_l^{l+1} \frac{1}{\sqrt{x}} dx = \gamma_1 \int_j^n \frac{1}{\sqrt{x}} dx = 2\gamma_1 (\sqrt{n} - \sqrt{j})$$

so that:

$$\sum_{j=n_0}^{n-1} \gamma_j \prod_{l=j}^{n-1} (1 - \alpha \gamma_l) \leq \gamma_1 e^{-2\alpha \gamma_1 \sqrt{n}} \sum_{j=n_0}^{n-1} \frac{e^{2\alpha \gamma_1 \sqrt{j}}}{\sqrt{j}}.$$

Checking that $x \mapsto \frac{1}{\sqrt{x}} e^{\alpha \gamma_1 \sqrt{x}}$ is non-decreasing on $[\frac{1}{(\alpha \gamma_1)^2}, \infty)$ it can be deduced that for any $j \geq n_0$,

$$\sum_{j=n_0}^{n-1} \frac{1}{\sqrt{j}} e^{2\alpha \gamma_1 \sqrt{j}} \leq \int_{n_0}^n \frac{1}{\sqrt{x}} e^{2\alpha \gamma_1 \sqrt{x}} dx \leq \frac{1}{\alpha \gamma_1}.$$

The lemma follows. □

References

- [1] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. JMLR: Workshop and Conference Proceedings vol 23 (2012), 23:39.1–39.26, 2012.
- [2] Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In Proceedings of the 22nd Annual Conference on Learning Theory (COLT), pages 2635–2686, 2009.
- [3] Jean-Yves Audibert and Sébastien Bubeck. Regret bounds and minimax policies under partial monitoring. Journal of Machine Learning Research, 11:2635–2686, 2010.
- [4] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research, 3:397422, 2002.
- [5] P. Auer, N. Cesa-Bianchi, Y. Freund, and R.E. Schapire. The non-stochastic multi-armed bandit problem. SIAM Journal on Computing, 32:4877, 2002.
- [6] Jean-Baptiste Bardet, Alejandra Christen, Arnaud Guillin, Florent Malrieu, and Pierre-André Zitt. Total variation estimates for the TCP process. Electronic Journal of Probability, 18:1–21, 2013.
- [7] Jean-Baptiste Bardet, Alejandra Christen, Arnaud Guillin, Florent Malrieu, and Pierre-André Zitt. Total variation estimates for the TCP process. Electron. J. Probab., 18:no. 10, 21, 2013.
- [8] F. Bouguet, F. Malrieu, F. Panloup, C. Poquet, and J. Reygner. Long time behavior of markov processes and beyond. ESAIM: Proceedings and Surveys, pages 193–211, 2015.
- [9] O. Cappé, A. Garivier, O. Maillard, R. Munos, and G. Stoltz. Kullback-leibler upper confidence bounds for optimal sequential allocation. Annals of Statistics, 41:1516–1 541, 2013.
- [10] Bertrand Cloez. Wasserstein decay of one dimensional jump-diffusions. preprint, 2012.
- [11] M. Duflo. Random Iterative Models. Classics in Mathematics. Springer-Verlag, New-York, 1997.
- [12] Stewart N. Ethier and Thomas G. Kurtz. Markov processes. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, 1986. Characterization and convergence.
- [13] H. Kushner and G. Yin. Stochastic approximation and recursive algorithms and applications, Second edition. Applications of Mathematics, 35. Stochastic Modelling and Applied Probability. Springer-Verlag, New-York, 2003.

- [14] H.J. Kushner and D.S. Clark. Stochastic approximation for constrained and unconstrained systems. Springer-Verlag, Berlin, 1978.
- [15] Damien Lamberton and Gilles Pagès. How fast is the bandit? Stochastic Analysis and Applications, 26:603–623, 2008.
- [16] Damien Lamberton and Gilles Pagès. A penalized bandit algorithm. Electronic Journal of Probability, 13:341–373, 2008.
- [17] Damien Lamberton, Gilles Pagès, and Pierre Tarrès. When can the two-armed bandit algorithm be trusted ? Annals of Applied Probability, 14:1424–1454, 2004.
- [18] K. S. Narendra and I. J. Shapiro. Use of stochastic automata for parameter self-optimization with multi-modal performance criteria. IEEE Trans. Syst. Sci. Cybern., 5:352–360, 1969.
- [19] M.F. Norman. On linear models with absorbing barriers. J. Math. Psych., 5:225–241, 1968.
- [20] H. Robbins. Some aspects of the sequential design of experiments. Bulletin of the American Mathematics Society, 58:527–535, 1952.